DOCUMENT RESUME

ED 344 569 IR 015 479

TITLE	Federal High Performance Computing and Communications Program. The Department of Energy Component.
INSTITUTION	Department of Energy, Washington D.C. Office of Energy Research.
REPORT NO	DOE/ER-0489P
PUB DATE	Jun 91
NOTE	51p.; Color photographs will not reproduce well.
AVAILABLE FROM	National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161.
PUB TYPE	Information Analyses (070) -- Reports - Descriptive (141)
EDRS PRICE	MF01/PC03 Plus Postage.
DESCRIPTORS	Agency Cooperation; *Computer Networks; Computer Science; Computer Software; *Computer System Design; Databases; Federal Programs; *Research and Development; Technological Advancement
IDENTIFIERS	*High Performance Computing; *Supercomputers

ABSTRACT

This report, profusely illustrated with color photographs and other graphics, elaborates on the Department of Energy (DOE) research program in High Performance Computing and Communications (HPCC). The DOE is one of seven agency programs within the Federal Research and Development Program working on HPCC. The DOE HPCC program emphasizes research in four areas: (1) HPCC Systems--evaluate advanced architectures for large-scale scientific and engineering applications; (2) Advanced Software Technology and Algorithms--research computational methods, algorithms, and tools; develop large-scale data handling techniques; establish HPCC research centers; and exploit extensive teaming among scientists and engineers, applied mathematicians, and computer scientists; (3) National Research and Education Network--research and develop high-speed computer networking technologies through a multiagency cooperative effort; and (4) Basic Research and Human Resources--encourage research partnerships between national laboratories, industry, and universities; support computational mathematics and computer science research; establish research and educational programs in computational science; and enhance K through 12 education. The stated goals of the DOE HPCC program are to support the economic competitiveness and productivity of the United States, accelerate the application of HPCC technology to the solution of scientific and engineering problems, and enhance U.S leadership in research, development, and deployment of HPCC technologies. A number of examples of DOE HPCC operations are provided, including the MatVu matrix visualization computer software package; Global Ocean Model; Molecular Dynamics Simulations; Pion Propagator; Three-Dimensional Tokamak Modeling; MediaView, a system for multimedia communication; and the CASA testbed, a wide-area, very high speed communication network that enables a number of collaborating agencies to work with geographically dispersed supercomputing resources. (DB)

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

DOE/ER-0489P

ED344569

The DOE Program Component of the

FEDERAL HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

US DEPARTMENT OF ENERGY
Office of Energy Research
Scientific Computing Staff
Washington, D.C. 20585

June 1991

BEST COPY AVAILABLE

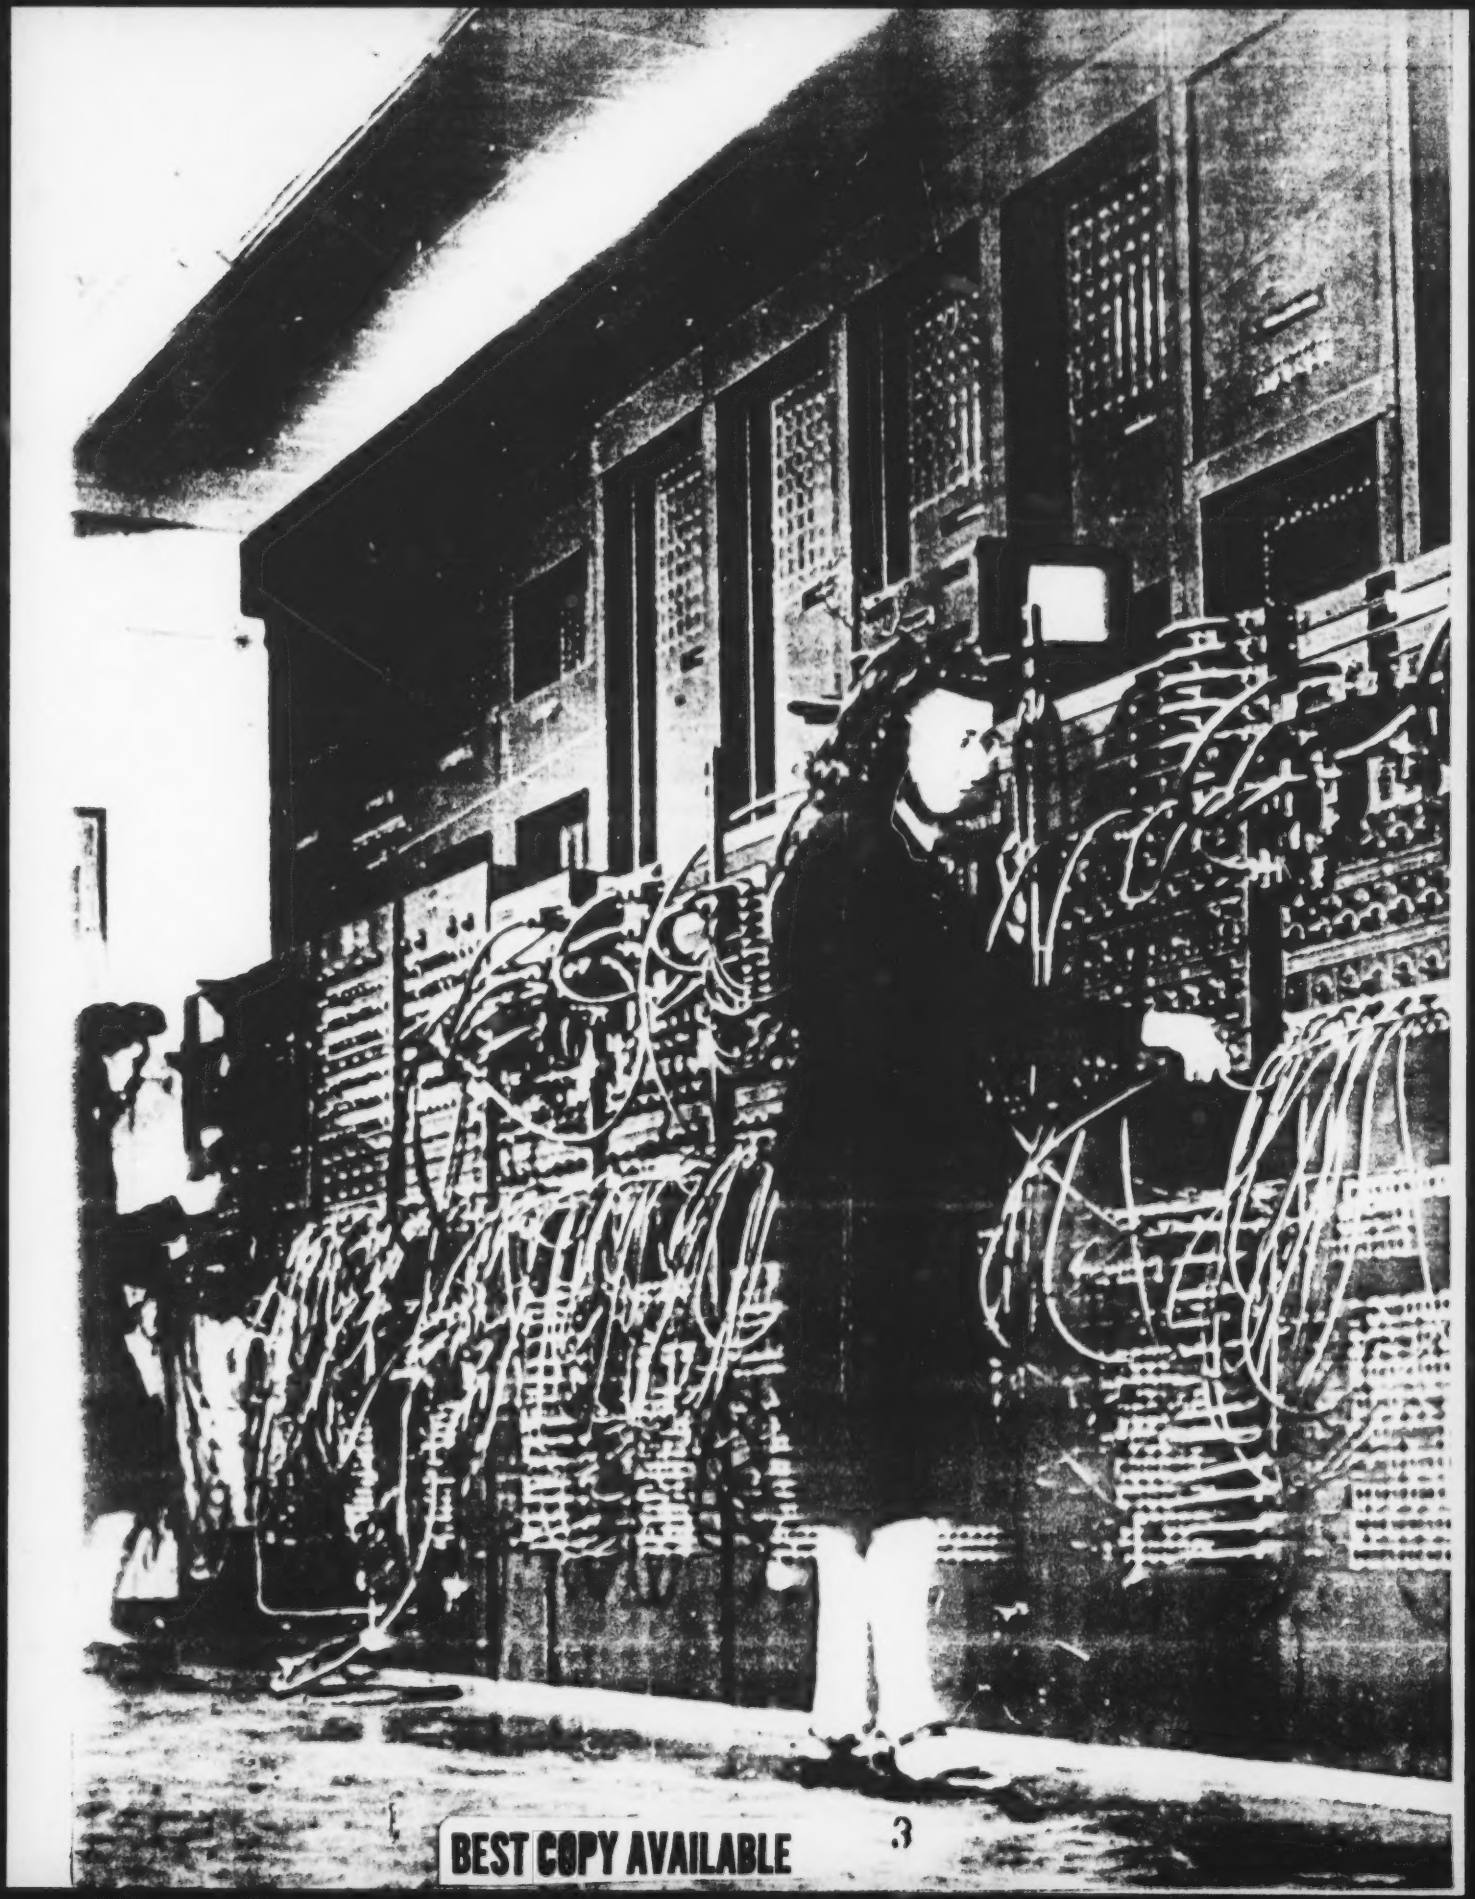

BEST COPY AVAILABLE

DOE/ER-0489P

The DOE Program Component of the

FEDERAL HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

US DEPARTMENT OF ENERGY
Office of Energy Research
Scientific Computing Staff
Washington, D.C. 20585
June 1991

ACRONYMS DEFINING SUPPORT AGENCIES

AFOSR	Air Force Office of Scientific Reseach
ARO	Army Research Office
AT&T	American Telephone and Telegraph
DARPA	Defense Advanced Research Projects Agency
DNA	Defense Nuclear Agency
DOE	Department of Energy
EPA	Environmental Protection Agency
IBM	International Business Machines
NASA	National Aeronautics and Space Administration
NSF	National Science Foundation

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831; prices available from (615) 576-8401, FTS 626-8401.

Available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161.

EXECUTIVE SUMMARY

The Department of Energy (DOE) High Performance Computing and Communications (HPCC) Program is one of seven agency programs within the Federal Research and Development Program in HPCC that was proposed with the President's Fiscal Year (FY) 1992 Budget and described in the supplemental report to the budget entitled *Grand Challenges: High Performance Computing and Communications.*[1]

The overall Federal HPCC Program is coordinated by the Federal Coordinating Council on Science, Engineering, and Technology (FCCSET) Committee on Physical, Mathematical, and Engineering Sciences (PMES) through its subcommittee on High Performance Computing, Communications, and Information Technologies (HPCCIT). The Grand Challenges report[1] describes the programs and interrelationships of the Federal agencies that are participating in the HPCC Program.

The purpose of this document is to elaborate on the DOE research program in HPCC.

DOE HPCC GOALS

To support United States economic competitiveness and productivity through interdisciplinary research and human resource development.

To accelerate the application of HPCC technology to the solution of scientific and engineering problems of significant departmental interest.

To enhance United States leadership in research, development, and deployment of HPCC technologies.

STRATEGY

Support computational advances through increased research and development efforts in areas of traditional strength in the department.

Promote the use of department and department-supported facilities as a market for HPCC prototypes and commercial products.

Support the underlying research, network, and computational infrastructures on which HPCC technology is based.

Develop the human resource base to meet the growing needs of industry, academia, and government in the area of HPCC.

PROGRAM DESCRIPTION

The DOE HPCC program will emphasize research in each of the following four key areas. Selected computational challenges, which have a significant effect on national leadership in science and technology, will be used as focal points for these efforts.

HPCC Systems—evaluate advanced architectures for large-scale scientific and engineering applications.

Advanced Software Technology and Algorithms—research computational methods, algorithms, and tools; develop large-scale data-handling techniques; establish HPCC research centers; exploit extensive teaming among scientists and engineers, applied mathematicians, and computer scientists.

National Research and Education Network—research and develop high-speed networking technologies through a multiagency effort.

Basic Research and Human Resources—encourage research partnerships between national laboratories, industry, and universities; support computational mathematics and computer science research; establish research and educational programs in computational science; enhance K through 12 education.

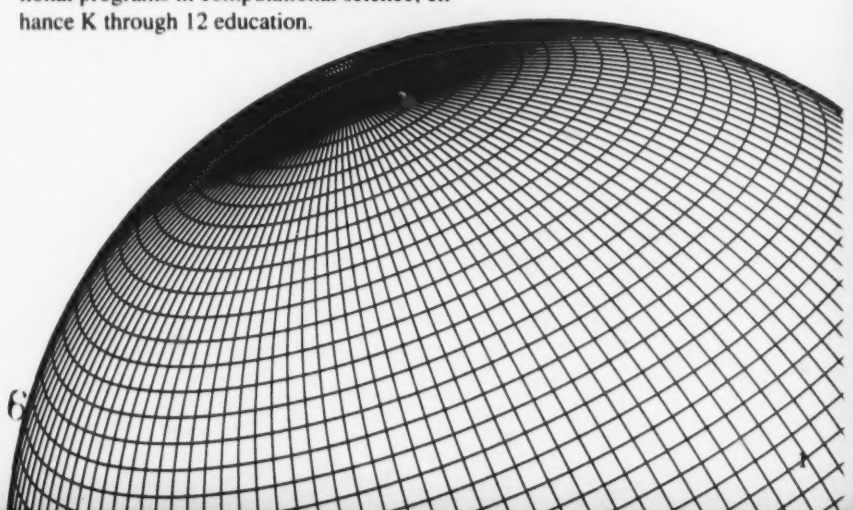

On February 5, 1991, Dr. Allan Bromley announced the FY 1992 U.S. Research and Development Program for high performance computing and communications to Congress. This HPCC Program, described in the supplement[1] to the President's FY 1992 budget, is the culmination of several years of effort on the part of senior scientists and managers in government, academia, and industry examining the state of U.S. high performance computer and network technology. The program recommends increased federal spending by the Department of Commerce, the Defense Advanced Research Projects Agency (DARPA), the DOE, the Environmental Protection Agency (EPA), the Department of Health and Human Resources, the National Aeronautics and Space Administration (NASA), and the National Science Foundation (NSF) for research in advanced computer technologies to develop dramatically more capable supercomputers, more powerful software capabilities, and high-speed computer networks. This DOE document describes and discusses only the potential DOE initiatives in response to the program plan and funding recommendations.

DOE HPCC PROGRAM HISTORY

Because of its mission and the computationally intensive nature of energy-related applications and problems, the DOE mission depends on advancements in computational techniques and computer and networking technologies. As a result, DOE has a long history of computational research and development, with strong industrial and university cooperation. The current DOE Applied Mathematical Sciences Program came about as the result of a suggestion by John von Neumann to enhance understanding of the use of digital computers in nuclear applications. Consequently, DOE has been prominent in maintaining the U.S. leadership in HPCC, in encouraging—and even providing—innovation in HPCC technologies, and in supporting US competitiveness and productivity through its extensive use of HPCC technologies. The table (right) outlines a history of DOE supercomputing hardware involvement in which the DOE national laboratory system has worked with U.S. vendors to bring promising and innovative computing technologies to bear on departmental applications.

Year / System / *Site*

1945 Eniac *U. of Pennsylvania, AEC(DOE)/LANL First User*

1952 Maniac *AEC(DOE)/LANL First Operational Fabrication of von Neumann Design*
1953 IBM 701 *AEC(DOE)/LANL, First customer installation*
1956 IBM 704 *AEC(DOE)/LANL, First customer installation*
1959 LARC *AEC(DOE)/LLNL, Prototype system*

1961 IBM 7030 Stretch *AEC(DOE)/LANL, Prototype; AEC(DOE)/LLNL, Serial 1*
1963 CDC 3600 *AEC(DOE)/LLNL, Serial 1*
1964 CDC 6600 *AEC(DOE)/LLNL, Serial 1*
1969 CDC 7600 *AEC(DOE)/LLNL, Serial 1*

1974 CDC Star 100
AEC(DOE)/LLNL, Serial 1
1976 MFENET
ERDA(DOE)/NMFECC at LLNL, First nationwide supercomputer network
1976 CRAY 1
ERDA(DOE)/LANL, Serial 1

1981 IBM 3081 *DOE/SLAC, First customer installation*
1982 CDC Cyber 205 *DOE/KAPL, Serial 3, first U.S. installation*
1984 CRAY 2/1 *DOE/NMFECC at LLNL, Serial 1 quadrant*
1985 CRAY 2/4 *DOE/NMFECC* at LLNL, *Serial 1 four processor*
1987 ETA-10E/4
DOE sponsored SCRI at FSU, Serial 1
1988 IBM
DOE/LANL first HIPPI prototype
1989 ETA-10G/4 *DOE/FSU, first and only 7-ns machine*
1989 CM-2 *DOE/LANL, first 64K floating point machine*

1990 CRAY 2/8 *DOE/NERSC, first 8-processor CRAY 2*
1990 Intel i860/128 *DOE/ORNL, Serial 1 Touchstone Gamma*

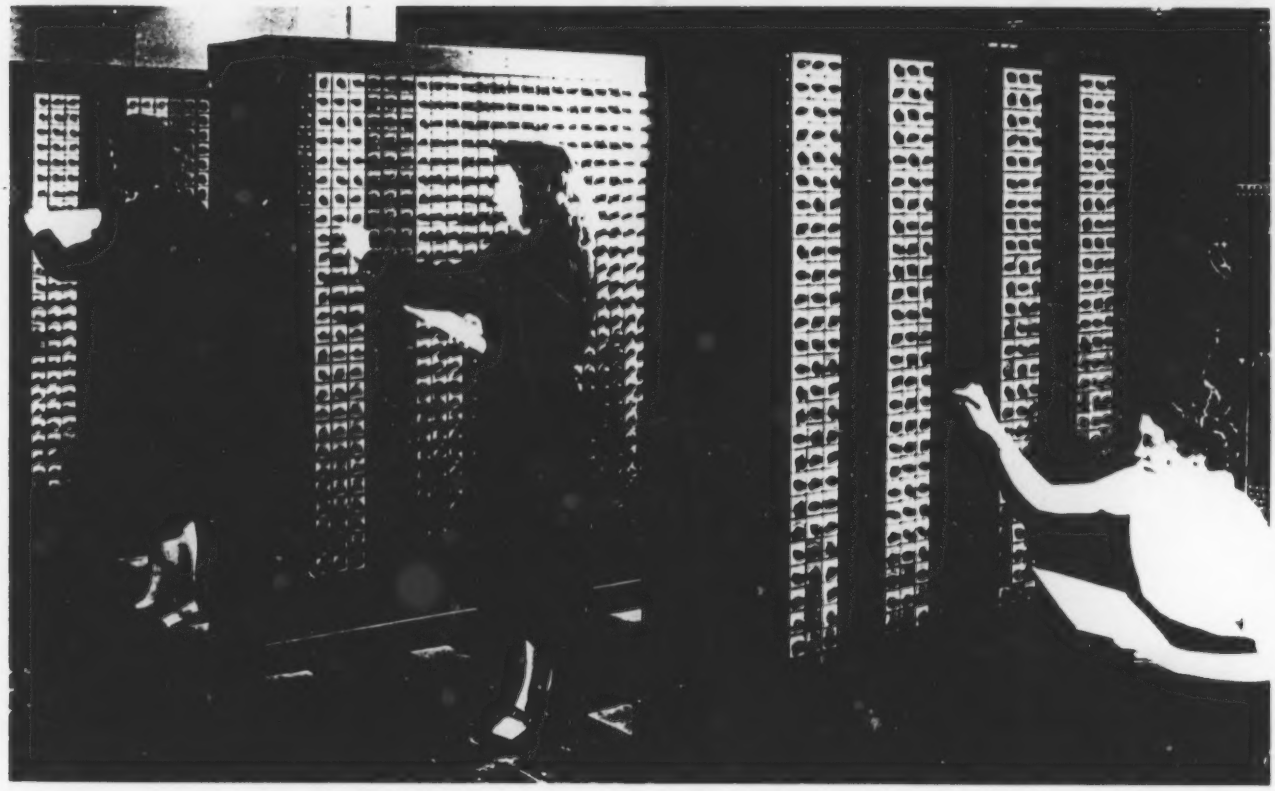

Eniac computer

DOE laboratories are among the pioneers of the network-based shared-resource architecture of the supercomputer environment and its supporting supercomputer system software. Examples include the first timesharing operating system for supercomputers, three generations of mass storage servers, a supercomputer Fortran environment including optimizing compilers, a portable Fortran mathematical subroutine library, interactive debugging, vector extensions, and a portable operating system interface library. The laboratories have also pioneered computational science and the supporting mathematical techniques and libraries.

The exploitation of state-of-the-art, innovative supercomputer technology has become more and more challenging. The complex, massively parallel computer architectures needed to provide the requisite computer power to solve forefront energy research problems in the next decade will present difficult computational research problems. These will require a rethinking of disciplinary and organizational boundaries and will demand interdisciplinary teamwork on a scale that is unprecedented in recent memory.

As an example, consider the computational challenges posed by the various energy systems and materials of the automobile. The DOE mission to provide and to implement a national energy strategy must include major efforts to reduce automotive gasoline consumption.

There are many energy-related problems—all of which are computational challenges—involved in developing more fuel efficient, environmentally sound, and safer automobiles. These include the modeling of the combustion systems, the materials and structure, the aerodynamics, the control of materials processing and components manufacture, and the use of alternative energy sources such as chemical batteries or solar technologies. Any one of these individual energy applications saturates supercomputer systems today. In order to develop and computationally experiment with the complex, integrated models required to advance the understanding of these problems, researchers will need to use the teraflops computer systems proposed for development in the Federal HPCC Program. It will also be necessary to bring together teams of scientists and engineers, ap-

John Von Neumann, circa 1950

BEST COPY AVAILABLE

plied mathematicians, and computer scientists to develop the software technology and algorithms for these difficult problems. High Performance Computing Research Centers (HPCRCs) will serve as focal points for these teams. In addition to addressing grand challenge applications problems, the activities of individual HPCRCs may include using early versions of advanced computers for software development, providing feedback to system developers, and providing network access to the broader research community.

The DOE has long recognized the need for computer networks to support remote access to its supercomputer systems and to support distributed research collaborations. In addition to supercomputers, the DOE is responsible for highly sophisticated facilities such as the Superconducting Super Collider, synchrotron light sources, neutron sources, and electron microscopes to probe the atomic and molecular structures for materials, medical, chemical, biological, and pharmacological research. High-speed communications access to such facilities is necessary to enhance the quality and scope of research collaborations. Over 15 years ago, the DOE implemented national computer networking access to supercomputers. And more than 25 years ago, DOE laboratories pioneered high-performance local network access and the client-server model of shared network services such as mass storage, I/O authentication, graphics, and terminal access. Because of this development and systems integration experience and because of its mission to operate these forefront facilities, DOE is in a unique position to contribute greatly to the advancement of U.S. science through the networking component of the HPCC Program.

The Federal HPCC Program includes a multiagency, multigigabit network research initiative. The DOE national laboratories have traditionally researched very high-speed networking technologies in cooperation with U.S. vendors to provide high-speed data transfer for supercomputers and for experiment control systems applications. Most recently, DOE has worked in collaboration with industrial partners to develop gigabit networks and interface technology that increase communications bandwidth by two orders of magnitude. Therefore, DOE also expects to contribute to the HPCC Program gigabit testbed projects—especially with regard to gigabit applications, protocols, and systems software.

In recent years, traditional theoretical and experimental techniques in science and engineering have been augmented by a powerful new technique: computational science. The term "computational science" is used to describe those intellectual activities in science, engineering, mathematics, and computer science that develop or exploit HPCC as an essential tool. While theory and experimental methods will not be replaced by computational science, they are being complemented by it in important ways. Currently, computational science educational programs are emerging in a small number of American universities. Having recognized the need for long-term investment in basic research and education in computational science, the DOE national laboratories have long been leaders in developing computational science techniques and in training limited numbers of postdoctoral researchers. DOE intends to contribute further in this vital area at all educational levels with a varied program, including teacher workshops in computational science, assistance with curriculum development, and access to high-performance computers for educational programs.

The DOE also recognizes that, in order to bring the full benefits of the HPCC Program to the national industrial complex and other entities concerned with proprietary or privacy information, proper security measures must be included in the program from the start. The DOE national laboratories have decades of experience in dealing with complex computing environments that process many levels of sensitive information.

The following sections provide detail on the research initiatives proposed by DOE as part of the Federal HPCC Program.

HIGH-PERFORMANCE COMPUTING SYSTEMS

The Federal HPCC Technology Development Program in High Performance Computing Systems will be primarily performed and coordinated by DARPA, as described in the Grand Challenge report.[1] The DOE will participate in technology development, primarily with regard to evaluating advanced architectures for large-scale scientific and engineering applications.

The DOE has long played an influential role in pushing the development of high performance computers. DOE applications are computationally and memory intensive. They also push the state of the technology for hardware interconnects to supporting communications, storage, and output devices. DOE has a natural interest in development of the foundations for the future generations of computing that will lead to advances in raw speed and capacity support for trillions of operations per second (teraflop) computing requirements. DOE also has an interest in encouraging the development of the next generation of machines so that architecture advances can be integrated into effective, usable computing environments. An important long-term goal that will guide DOE research in this area is to foster development of high performance computational facilities that are architecturally balanced in a way to make them usable by general DOE applications.

DOE is specifically interested in research that will lead to understanding the architectural limitations that will shape future machines, understanding how basic parallel software support is impacted as the architecture scales to higher performance, and providing resources for research in developing basic components of future computing environments.

MASSIVELY PARALLEL COMPUTERS

Future generations of machines are expected to rely heavily on architectural parallelism. Identifying and understanding the limiting factors in parallel computing equipment is essential in establishing realistic upper bounds on the requirements being placed on machine design and supporting technologies. Computational models that require true teraflop solution speeds will saturate the underlying hardware in many areas. We need to understand where these saturation points occur and seek to establish the means to overcome them. The DOE will conduct research to determine the effects of increasing parallelism in both hardware and software. Among the many issues to be addressed are the impact of code size on computational requirements, the effect of various languages and models of computation on the amount of usable parallelism made available, the impact on reliability from increasing hardware complexity, the need for operating system features that optimize the use of the hardware, and the increased demands placed on supporting technologies.

Massively parallel computers use an approach to achieving speed that is radically different from that used in today's production supercomputers. In massively parallel computers, many hundreds or thousands of simpler and less expensive processors are coupled with a large amount of memory. These processors team up to carry out computations in parallel by breaking the problem into many subtasks that can be carried out simultaneously by the individual processors. The processors exchange information as needed to complete the computation. In this way computations can be carried out many times faster on the massively parallel computer than on one of its individual processors.

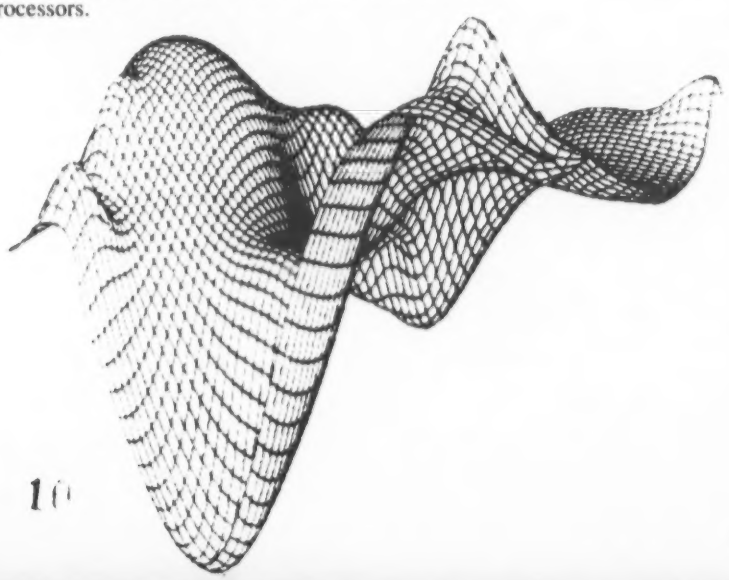

There are many different ways to link up processors and memory in massively parallel computers. This adds to the flexibility of the computer but also complicates the job of computer design: system architects are faced with many alternatives and choosing the "best" ones can be a difficult task and may be strongly application dependent.

The current generation of massively parallel computers has been successfully used to solve selected applications problems as fast as or faster than the best vector supercomputers. Because of improvements in manufacturing technology, it is forecast that in two years these machines will calculate ten times faster than they do now; in five years they may be hundreds of times faster.

To see why this speed is needed, consider an example of three-dimensional flow simulation used for global climate or reservoir modeling. Current models, in which the flows are calculated by repeatedly evaluating equations millions of times, may not be sufficiently accurate. Nevertheless, these calculations can take hundreds of hours on a conventional supercomputer. Accurate climate models could require solutions of equations hundreds of millions of times at each time step of the simulation. Completing such calculations would require years of supercomputer time. For weather forecasting, we need the data in hours. Massively parallel computers currently being designed could do these demanding calculations in days; massively parallel computers proposed to be developed under the HPCC Program could meet this need.

GRAPHICS AND VISUALIZATION

The volume of data produced by such powerful computing machines will tax storage media, input/output devices, and the networks that carry the data to the user. The only reasonable way to review the data is visually. For example, accurate three-dimensional flowfield simulations might produce as many as ten billion potentially important bytes of data at each timestep. A typical simulation would involve thousands of timesteps. Thus, there might be as much as ten trillion bytes of data to be processed from one problem. If this document had that much information in it, it would be a few billion pages long! Of course, only a very small fraction of that data could be used in any human endeavor. Visualization enables us to represent information in a form that effectively uses human perceptive capabilities.

NETWORKING

To be effective, HPCC must be available to a large user population, not only in the national laboratories but in universities and industry as well. A goal of the HPCC Program is to disseminate this computing capability throughout American science, technology, commerce, and industry. This will require local and national networks with capabilities matched to the computing resources being accessed. Further, within a computing center, the networks to tie machines together with one another and with storage devices will need awesome bandwidths—billions of bytes of information per second or more. The technology to carry out this networking function is part and parcel of the same technology in the computers that generate the data: fast switches, high-speed interconnects, networking protocols and services, and increasingly sophisticated network-server computers.

WORKSTATIONS AND PERSONAL COMPUTERS

The impact of massively parallel supercomputers will be complemented by new generations of personal workstations with unprecedented power. This power will spill over into the personal computer (PC) market. The cost of low-end workstations is now comparable to that of well-equipped PCs. PCs are now acquiring workstation-like functions. This trend will continue until there is little to distinguish the two classes. It should not be surprising that workstation and PC will share in the power revolution: many of today's massively parallel computers use processors that grew out of PC and workstation processor technology.

Today's personal workstations provide the power of yesterday's mini-supercomputer for a few thousand dollars—a major achievement

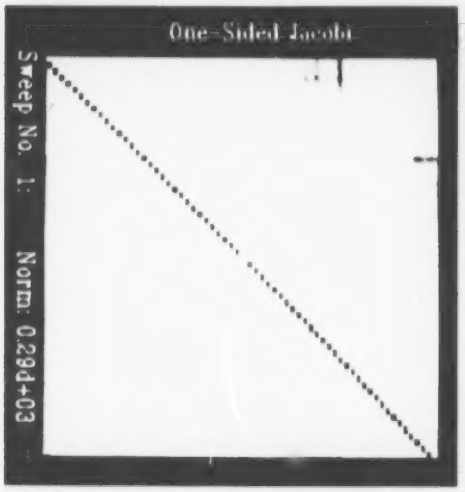

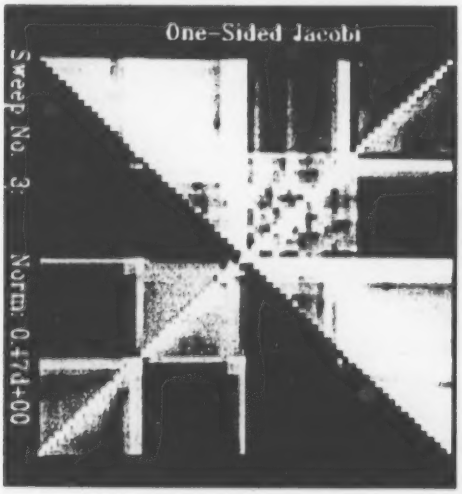

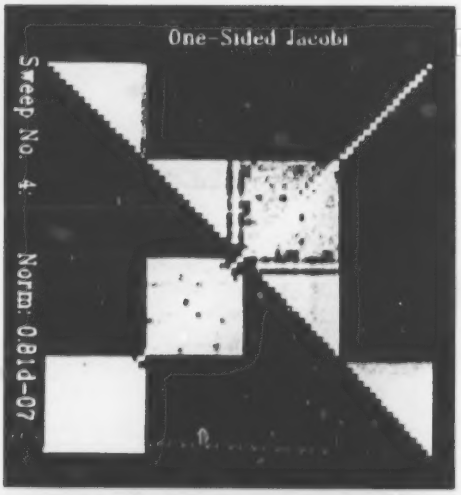

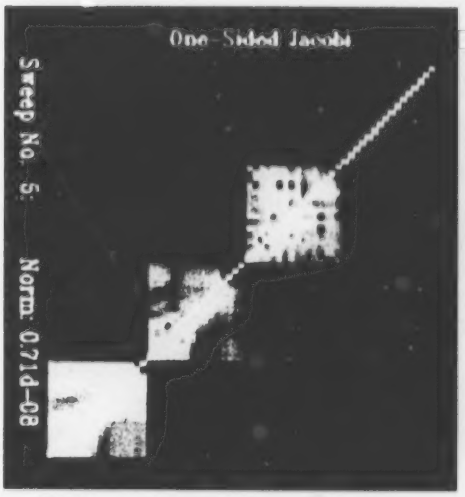

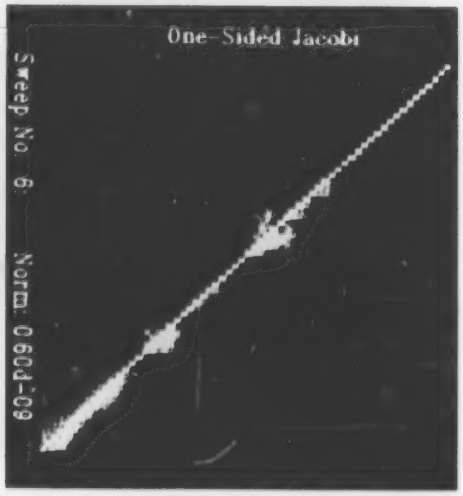

ith the use of MatVu, a matrix visualization package, researchers at the Center for Supercomputing Research and Development have been able to classify the convergence patterns of numerous eigenvalue and singular value spectra by assigning a logarithmically scaled color table to the magnitudes of the off-diagonal elements in each sweep of the particular Jacobi algorithm.

Parallel Jacobi methods (two- and one-sided) on the Alliant FX/8 and CRAY X-MP/48 have been quite effective for computing the eigenvalues and eigenvectors of rectangular matrices.

The color table used in the images ramps from black (matrix elements of magnitudes less than or equal to the 64-bit machine precision) through blue, green, yellow, and orange to red (matrix elements of largest magnitude). The diagonal elements (yellow, orange, and red) in each of the matrices represent improved approximations to the exact singular values, with S_0 yielding the most accurate approximation.

Work supported by DOE, NSF, AFOSR, AT&T, and IBM.

AVAILABLE

of American industry. In two years workstations may have the power of a CRAY Y-MP processor on a variety of applications. Much of the work currently suitable only for supercomputers will be feasible on inexpensive workstations.

COMPUTING ENVIRONMENTS

An exciting aspect of the HPCC Program is its promise—at relatively modest investment—to revolutionize our approach to science, commerce, and industry. The combination of massively parallel supercomputers, data management and visualization computers, ultra-fast networks, and very powerful personal workstations will greatly increase our knowledge base and speed the rate of scientific progress. Similarly, the design process in industry will be enhanced: rapid prototyping will be a reality, the cost of complex designs should decrease, and product quality should benefit greatly. Accessing immense commercial databases will become an integral step in business planning. To the extent that information and knowledge are the currency of tomorrow, the HPCC Program will provide a differentiating technology for America.

DOE's challenge is to help foster continued development of the diverse pieces of the HPCC Program, but especially to provide the expertise to glue the pieces into a seamless whole.

HIGH PERFORMANCE DATA STORAGE SYSTEMS

High performance data storage systems need to be developed that are able to effectively serve an environment of massively parallel computers, general-purpose supercomputers, scientific workstations, and visualization.

One of the very important and very difficult challenges of this decade will be to satisfy the data storage and data access requirements in the HPCC environment. Current large problems that run on supercomputers generate from one to ten gigabytes of data. This data must be saved and then be available for quick access. With the current success in moving

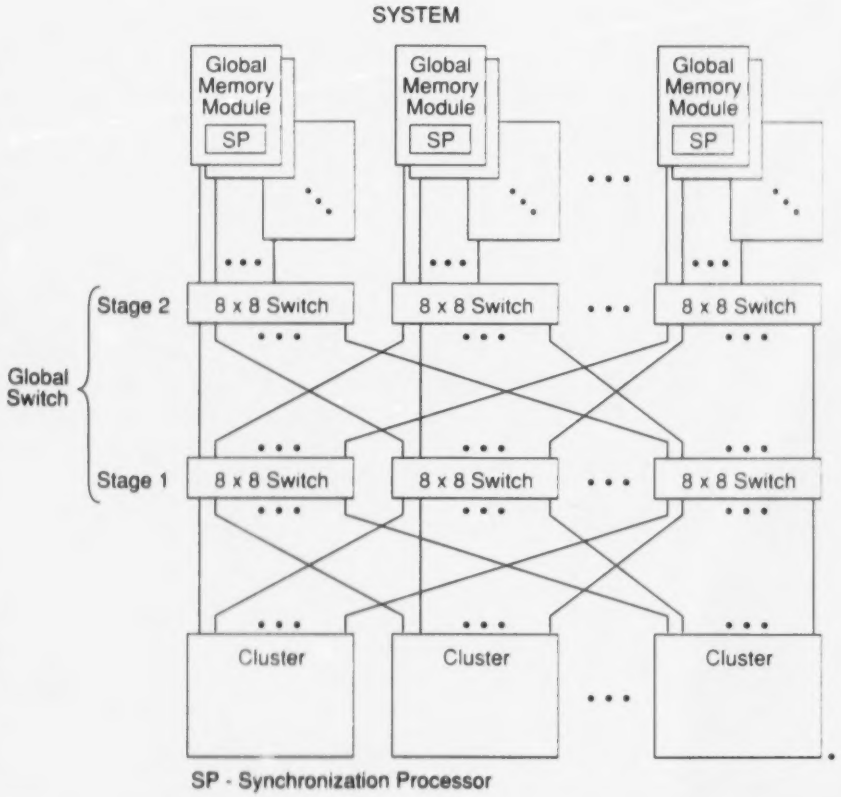

edar, the experimental shared-memory multiprocessor system of the Center for Supercomputing Research and Development, has a scalable hierarchical structure. Clusters, each composed of eight processors, are connected to a global memory by means of a very high bandwidth shuffle-exchange global network. The shuffle-exchange network is architecturally scalable because it has a fixed number of lines per I/O node, regardless of network size. In addition, the system exhibits scalability as a result of the power of the network and the hierarchical memory provided by the clusters. Cedar exhibits stable performance over a wide range of applications because of this two-level parallelism, with and between clusters, for fine and coarse grain parallelism, respectively.

Work supported by DOE, NSF, AFOSR, and IBM.

The Advanced Computing Research Facility (ACRF) at Argonne National Laboratory is committed to research on computers with innovative designs, with parallel architectures as the principal focus of attention. The ACRF comprises a wide variety of advanced computers, ranging from a four-processor graphic supercomputer to a 16,000-processor massively parallel Connection Machine.

This diversity of machines enables researchers to conduct experiments on software portability among different computer architectures and to evaluate the suitability of different configurations, including shared memory and distributed memory, for specific applications. The machines are used in a wide range of research projects, including design of parallel algorithms, analysis of languages for programming multiprocessor systems, and development of parallel programming methodologies for writing parallel programs.

Established in 1984, the ACRF has become recognized as one of the world's leading centers for parallel computing research. Approximately three hundred researchers from industry, universities, and other government laboratories use the ACRF machines each month for studies in software and algorithm development.

To encourage use of the facility, the ACRF sponsors classes, workshops, and institutes on parallel computing. More than 500 scientists have participated in these activities, and many of these researchers continue to use the ACRF machines remotely via national networks. The ACRF also has established two affiliates programs to promote the transfer of research results to industry and academia. Currently, 15 industrial affiliates and 27 university affiliates are using the ACRF advanced computers to conduct computational research at the leading edge of technology.

Work supported by DOE and NSF.

large problems to massively parallel computers such as the Connection Machine, the data storage and data access requirements have dramatically increased. A large problem on the Connection Machine will generate from tens of gigabytes up to a terabyte. To fully support the current state-of-the-art, large memory, massively parallel supercomputer, it would be necessary to achieve data transfer rates of at least 50 megabytes/second and a storage capacity of at least a hundred terabytes of data. As the massively parallel machines become more powerful, the data handling requirements will likewise increase.

To meet these data storage and data access requirements, it will be necessary to develop very high performance storage systems with advanced data management capabilities. Such storage systems must be scalable to meet the increasing requirements and to maintain a balanced overall HPCC environment.

ADDITIONAL DOE ROLES IN HIGH PERFORMANCE COMPUTING

DOE continues, through its leading role in the computational sciences, to have a strong interest in shaping the development of future HPCC environments. An important goal for DOE is to foster architecturally balanced approaches usable by a broad class of applications. These applications are the foundation for the basic research needs of many U.S. industries.

Specific DOE interests include (1) issues associated with the scalability of architectures and their impact on software, (2) research aimed at providing the resources needed for components of future systems, (3) reliability, and (4) integration of the components of this distributed computing environment.

The DOE laboratories have a great deal of expertise in simulating advanced computer architectures. This capability will be important to investigations of the scalability of multiprocessor architectures. It can also enable accurate prediction of the performance of important applications on new designs before they reach hardware production. Thus, design flaws can be remedied early in the product development cycle before changes become prohibitively expensive.

ADVANCED SOFTWARE TECHNOLOGY AND ALGORITHMS

Software and algorithms are the keys to making HPCC systems both useful and economically successful. Furthermore, for earlier generations of computers, there is clear evidence that improvements in software and algorithms lead to larger performance gains than did improvements in machine architecture. However, the computing systems of the past had relatively simple architectures, making it possible to largely decouple applications algorithm and software development from the details of machine architecture. Today, the simple architectures have been thoroughly exploited. The complex architectures necessary to carry HPCC through the present decade require extensive development of new software technology, numerical methods, and applications algorithms in order that their potential be realized.

The Center for Computational Engineering (CCE) at Sandia National Laboratories, Livermore was established to develop and apply computational technology to problems of national interest. Sponsored by the Department of Energy and private industry, the CCE is tasked with bringing together researchers and computer scientists from government and private industry to conduct mutually beneficial scientific research.

Central to the CCE is the new generation of massively parallel supercomputers. To be able to use these massively parallel supercomputers, existing software must be rewritten in parallel form and new numerical methods must be developed. Initial work will be with sponsors in four separate modules: global climate change, macromolecular design and environmental health software engineering, and field data management.

Through the CCE's modules, Sandia will help corporate sponsors apply the new supercomputers to problems in their own areas of expertise. In other words, the modular structure will facilitate technology collaboration. By pooling the talents of scientists from industry, universities, and national laboratories, the CCE will automatically transfer computational technology, thereby helping sponsors reduce the concept-to-design cycle and speed product introduction. The sponsors' industrial expertise will help focus the CCE's research on real-world applications.

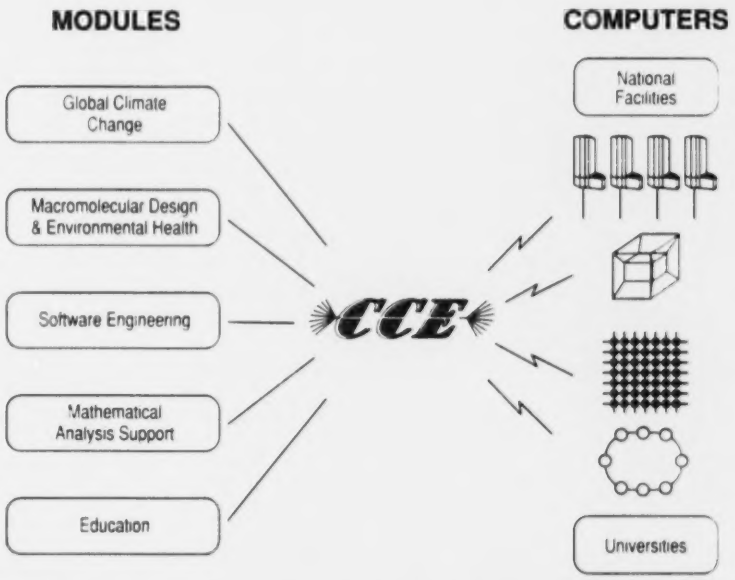

The CCE links engineering "modules" to evolving parallel processing technology.

Work supported by DOE.

Addressing this challenge requires a rethinking of disciplinary and organizational boundaries and an extraordinary teaming effort. Integrated approaches, reaching all the way from computer hardware design to grand challenge applications, must be sought. In order to create the software and algorithms for parallel, distributed, and hierarchical computing, computer scientists and applied mathematicians will need to better communicate with scientists in these important applications areas.

Most of the Federal HPCC generic software technology development will be funded by DARPA. Because of the critical need for interagency coodination in this area, all participating agencies will coordinate their advanced software technology and algorithms (ASTA) programs through the HPCCIT working group.

SUPPORT FOR GRAND CHALLENGES

The DOE, both as a natural consequence of its mission and specifically by virtue of its need to solve many of the grand challenge problems, has at its disposal a broad range of important applications on which to base a strong and aggressive software technology and algorithms support effort. To be successful, this effort will need to operate in close collaboration with grand challenge researchers so that a detailed understanding of applications will be incorporated into the earliest stages of software and algorithm design and so that application code design decisions will be made in light of full understanding of feasible software and algorithm options. Grand challenge research teams will be specifically called on to test emerging software and algorithms and to recommend modifications and improvements. Early transfer of developed technology to the private sector will be vigorously pursued to provide both financial leverage and additional feedback on the effectiveness of the DOE-supported development efforts.

Energy Conservation and Fossil Fuel Combustion

For the foreseeable future, 90% of the energy needs of the United States will be met by the combustion of fossil fuels. Fossil fuels are burned in stationary combustion chambers, for example, for electrical power generation, and in mobile forms such as in automobiles.

Automobile engines are most efficient when run at high temperatures (Carnot effect), but increased temperature leads to increased nitrogen oxide emissions. The burning of alternative fu-

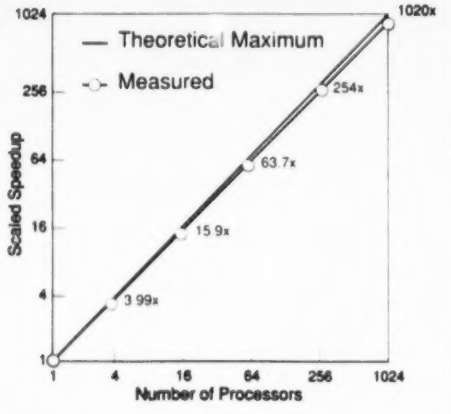

ioneering research at DOE's Sandia National Laboratories in parallel methods, applications, and performance evaluation on a 1024-processor NCUBE/ten hypercube won the Karp Challenge and the inaugural Gordon Bell Award in 1988, as well as an R&D 100 Award in 1989. This research demonstrated 1000-fold parallel speedups for three full-scale scientific applications on a commercial parallel processing system and introduced three new concepts of performance modelling: scaled speedup, fixed-time speedup, and operation efficiency. Sandia was given a second R&D 100 Award in 1989 for the Linda Project, its innovative research in parallel distributed computing using a large number of networked heterogeneous computers.

Work supported by DOE.

els such as methanol is complicated by the emission of formaldehyde, a known carcinogen, into the atmosphere. Pollutants are affected by local geology and climatic conditions, making it necessary when seeking solutions to take into account the total system of fuel, engine, and atmosphere. Our environment is too delicate to be used as a testbed; therefore, we must use supercomputers to simulate the atmospheric effects before experimenting.

An important example of the application of DOE supercomputer methodology is in the development of an engine design code that is fully three dimensional and includes models of chemical reactions, fluids, gases, and particulate formation. The code is in a constant state of evolution and is designed to handle the most complex engines, such as the stratified charge engine and the two-stroke engine.

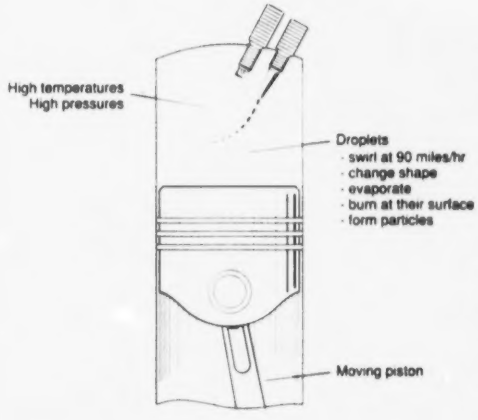

detailed understanding of fuel burning inside an internal combustion engine requires the comprehension of a number of complex processes. One example is the Direct-Injection Stratified Charge (DISC) engine, which requires high performance computing for its optimum design. DISC is an experimental engine design. Its goal is to run leaner and cooler than conventional engines, thereby reducing emissions while increasing the compression ratio and control of ignition for greater efficiency and fuel economy. The burning of the fuel-air mixture in such an engine brings into play a few hundred different chemical reactions between numerous short- and long-lived chemical species. The rates at which these reactions occur depend on the temperature and concentrations of the species. For instance, as combustion proceeds, various pollutants may be formed, such as nitrogen oxides (NO_x) and unburned hydrocarbons. These hydrocarbons may form solid particles known as soot. Also, the fuel may autoignite in some region before the flame arrives, giving rise to the small explosion known as engine knock.

Because engine operation is so complex, designing an optimum engine by experimentation (varying all design parameters independently) is prohibitively expensive. A computer model of the engine allows us to effect design changes numerically on a scale of hours at a fraction of the cost.

The engine components depicted in the illustration are a cupped piston, fuel injector, and spark plug. The piston may travel at speeds of up to 90 miles per hour. It draws in fresh air and expels exhaust gases as the combustion cycle progresses. The air and exhaust gases flow at approximately the piston velocity. At an appropriate point in the cycle (upper right in the figure), liquid fuel is injected into the engine cylinder. The fuel jet breaks up into small droplets, which interact with the air. The droplets alternately elongate and flatten; some of them break up into two new droplets, others collide and coalesce to form larger drops. The spray is caught up by the air flow and swirls around in the cylinder. It simultaneously evaporates to form a gaseous fuel/air mixture.

Shown on the right is a computer simulation depicting four different times during the combustion cycle of a DISC engine. The multicolored surface indicates the fuel's location. When the spray has evaporated sufficiently a spark ignites the mixture, and combustion begins as a flame propagating through the chamber. As the engine proceeds through its cycle (indicated by the increase in crank angle from –23 to +17 degrees), the fuel burns, resulting in high temperatures (see color code).

Each of the models within the code represents an approximation to the facts, sometimes because the facts are not well understood but often because the capabilities of existing computers do not allow such information to be included. For example, the over 400 chemical rate processes involving hydrocarbon and nitrogen chemistry are treated globally by less than ten such reactions in order that the code run in a few hours on a large supercomputer. Since, however, the 400 reactions are known from first principles, the problem is addressable: what are required are better algorithms running on a machine 10,000 times more powerful, a teraflop machine.

This computational design technology is used by many private industrial engine design firms as well as universities and government laboratories.

Present engine models are deficient in one way or another because of limitations in computing resources. Computer codes, such as KIVA, can provide information on large scale fluid flows, but are limited in their ability to describe the small scale flow structures in turbulent flow. Furthermore, fuel chemistry models are necessarily crude. Current chemical kinetics calculations approach the detail necessary for an adequate description of the chemistry, but even these models require serious improvements. Such calculations are capable of only the most rudimentary fluid mechanical effects.

The envisioned increase in computing power is critical to the effective development of clean, efficient engines.

Work supported by DOE.

BEST COPY AVAILABLE

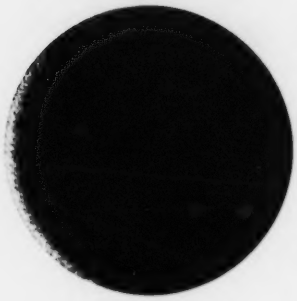

global ocean model running on the Los Alamos CM-2 is the first realistic general circulation model to be implemented on a massively parallel machine. The code was rewritten for the CM-2 based on the Cray version of the Semtner-Chervin global ocean model, which is highly parallelized to run on the multiprocessor CRAY X-MP or Y-MP. For large problems, the CM-2 code runs with speeds comparable to a full 8-processor Y-MP. New algorithms were implemented to substantially improve the code's performance. The model is a finite difference code with realistic surface and ocean-bottom topography.

To fully resolve the complicated currents and circulation in the ocean requires grids with spacing substantially less than 50 kilometers. The largest calculations that have been run to date with the Cray version used grids with 55-kilometer resolution. With the current CM-2, it is possible to solve problems with a 28-kilometer resolution. The ultimate aim in global climate studies is to develop an advanced climate model capable of describing the fully coupled ocean/atmosphere system, which will be necessary for solving important problems, such as predicting the effect of greenhouse gases on global warming.

The image shows the observed temperature (at 160 meters depth), which is used in model calculations. The warmest areas, in the equatorial region, appear red; the coolest areas, near the poles, appear orange and yellow; and the areas in between are shown in blue and green.

Work supported by DOE.

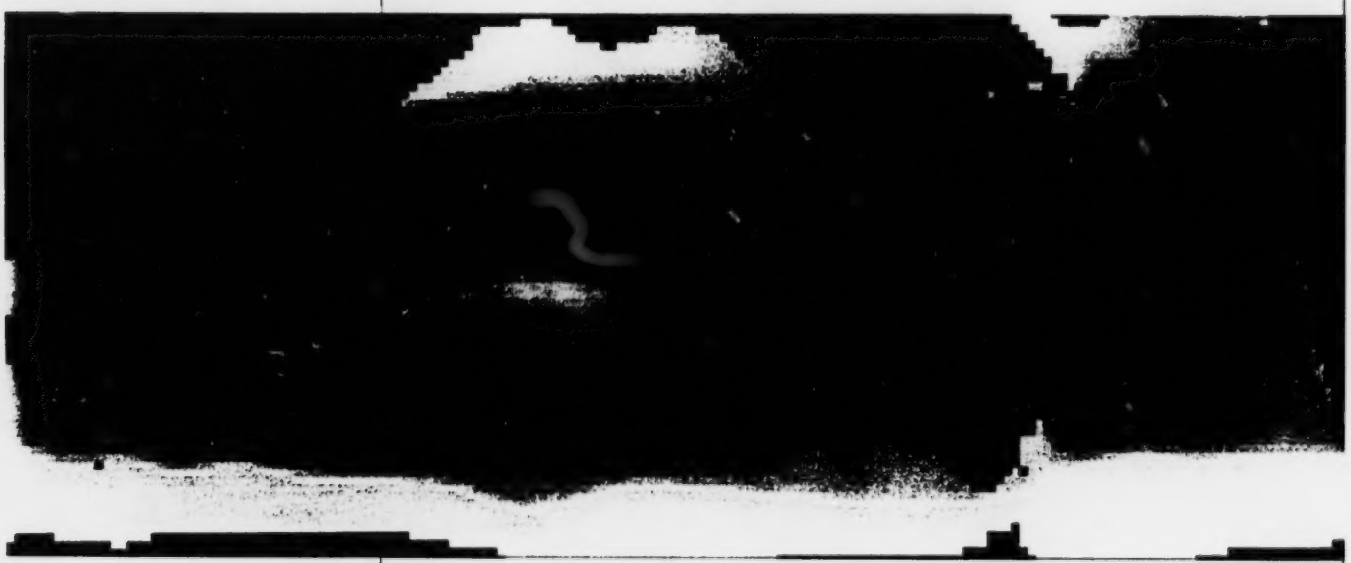

Global Climate Modeling

It is paramount that we improve our understanding of global climate and its potential impact upon human activities and the environment. Generation of carbon dioxide, methane, and other greenhouse gases as a by-product of energy production and other activities must be better quantified, and our capability to predict their influence on climate must be greatly enhanced if we are to formulate rational energy strategies and policies. Numerical computer models provide the only means for projecting the impact of greenhouse gases on future climate. Present-day global climate models are able to simulate satisfactorily some aspects of the current climate, but comparison of future climates predicted by various models reveals significant and disconcerting disagreements.

A new generation of models that have finer spatial resolution and more realistic treatment of the critical physical and chemical processes that control our climate is needed. In addition, longer computations spanning many decades of simulated time (200 years or more) will be needed. Current models, however, are constrained by the resources of even the largest

supercomputers now available. CHAMMP—the Computer Hardware, Advanced Mathematics, and Model Physics Climate Modeling Program—is a DOE program specifically inteded to develop within 10 years advanced climate models with the above-mentioned improvements and also capable of much longer simulations.[2] The advanced computer systems and software technologies being developed by the DOE HPCC Program will provide capabilities critical to the success of CHAMMP and to solving the nation's environmental and energy-related problems.

Biosciences

Biological research plays a major role in improving the quality of our lives. Computer use is, currently, a major activity of structural biology research and will become increasingly important. There are two categories of use: (1) data acquisition and analysis and (2) modeling and theory. Also noteworthy are the efforts being made to map and sequence the human genome—the full set of instructions for a human being. Such research is crucial if we are to determine the genetic basis of many human diseases.

Currently one of the principal uses of computers is for data acquisition and analysis. Efficient use of high-intensity photon and neutron beams provided by existing facilities demands online computer control of beams, spectrometers, and detectors. Data acquisition rates from diffraction and nuclear magnetic resonance (NMR) experiments are high, and analysis of data using established algorithms, an essential component of most structural biology research sponsored by the DOE Office of Health and Environmental Research, is very computer intensive. Current computer resources are not adequate to support the efficient use of existing facilities. Modeling of macromolecular dynamics for crystallographic data refinement and electrostatic field calculations for determination of structures in aqueous solution are extremely demanding of computer power and time. These types of simulations are important in that they can provide details of processes that cannot be easily probed by experimental techniques, such as the relaxation of water molecules in the vicinity of biological molecules.

In the future, efficient use of high flux beams at High Flux Beam Reactor and National Synchrotron Light Source and the advanced capabilities of the Manuel Lujan Jr. Neutron Scattering Center, the Advanced Light Source, the Advanced Photon Source, and the Advanced Neutron Source will require application of advanced and increasingly computer-intensive data acquisition and analysis protocols. Requirements for computer-intensive molecular dynamics and electrostatic calculations will increase as more structural biology data is acquired and analyzed. Development of parallel processing and computer graphics will be required to support online structural analysis and computer-based search strategies for sequence/structure or structure/function correlations.

The human genome initiative requires algorithms for comparing the "fingerprints" of pieces of DNA that are known to be derived from a larger piece of DNA but whose degree of overlap and spatial ordering relative to one another is not known. Determining these relationships is highly computationally intensive and stretches the art of combinatorics. Moreover, the requirements for this work are far from static: the quantity of sequence data has been doubling every year or two. If the human genome initiative is successful, the quantity of data 15 years from now will be 100 to 1000 times the current amount. The sequencing effort will require computational power, unprecedented in the realm of molecular biology, for improved data management and analysis, and communications links between those generating, managing, and accessing the experimental and interpretive data.

Materials Sciences

Many of the problems associated with materials research require computation of the subtle many-body effects that manifest themselves in high-temperature superconducting materials and the properties of complex polymers. Challenges in the realm of computational physics, chemistry, and engineering include (1) elucidating the mechanisms of chemical catalysts, the design of new catalysts, and the design of new materials for separations science; (2) understanding the electronic properties of novel materials, such as high-temperature superconductors, polymers and synthetic metals; (3) extending the empirical molecular mechanics approaches currently employed in the pharmaceutical industry for the design of drugs; and (4) designing new materials and alloys with novel properties based on an understanding of the interaction among the material constituents at the molecular level. Initial small applications of these challenges tax the capability of today's most powerful computing technologies and will help to influence their future directions.

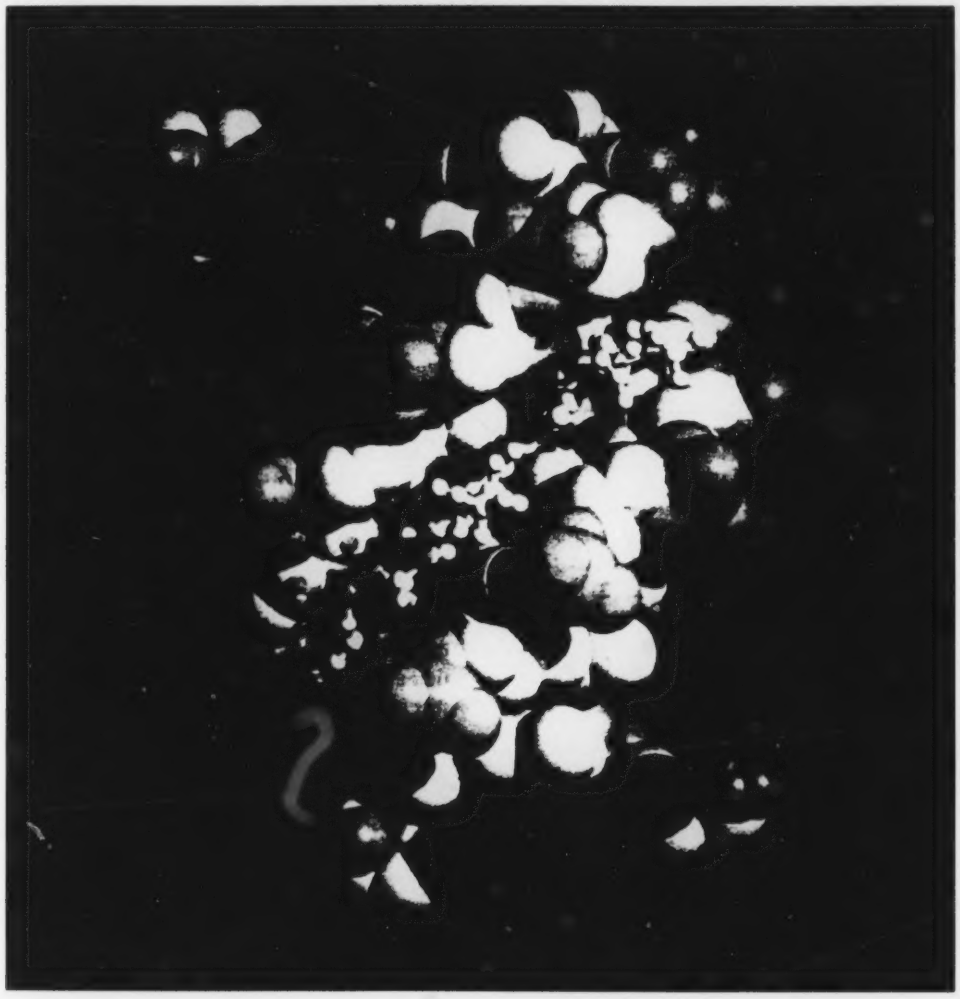

While experimental techniques such as two-dimensional NMR spectroscopy can provide detailed information about the relative distances between atoms in biological molecules, the translation of this information into three-dimensional structural information can be enhanced considerably by the use of molecular simulations. A recent study of interactions in drug-DNA complexes, performed at the Los Alamos National Laboratory, was carried out to construct a set of models consistent with two-dimensional NMR data. The particular case involved a complex interaction shown in the figure between a specific anticancer drug (Distamycin-2) and a segment of DNA. The atoms of the anticancer drug are represented by the small cyan objects. The surface of the DNA is indicated by larger atoms, where different colors represent different bases: blue is guanine, green is cytosine, yellow is adenine, purple is thymine, red represents the charged phosphate groups, and white represents sugar.

A molecular dynamics simulation followed by energy minimization showed that there are actually three distinct structures consistent with the NMR data and that in two of these structures the DNA molecule is bent significantly. In a nonlinear mode of motion shown by the drug-DNA complex, each of the three distinct minima is sampled every 5 ps (10^{-12} s). The combination of two-dimensional NMR and computer simulations gives a much clearer interpretation of the structure and dynamics of drug-DNA complexes than would be obtained by simply attempting to fit the experimental data. Calculations of this type can also provide guidance for future studies of protein-DNA interactions at the atomic level.

Work supported by DOE.

key aspect of designing improved materials is the understanding of their response to various kinds of dynamic loading, including fracture properties (such as the yield stress and the mode of fracture). Given a model for the interatomic forces, the fracture process can be simulated using molecular dynamics (MD) techniques, provided that the trajectories for a macroscopically large sample of atoms can be followed for a sufficiently long time on a computer. Simulations of ≈10^{12} atoms (i.e., a block 1 μm on edge) are required to properly model the effects of grain boundaries, plastic deformation, and crack propagation. Until very recently, the largest MD simulations have involved 10^4 to 10^5 atoms, so that no realistic calculations of fracture in three-dimensional solids have yet been attempted.

An exceptionally fast MD computer code is being developed to exploit massively parallel architecture of the Connection Machine (CM-2). It will perform simulations on >10^8 atoms, corresponding to a two-dimensional square of material approximately 0.25 μm on edge. Many of the features of three-dimensional fracture are retained in two-dimensions, so that these simulations will make the first direct connections between the atomic interactions and the macroscopic fracture properties.

Shown here is a sequence of snapshots from a sample run of a spall process, using ≈30 K atoms, interacting with a many-body potential appropriate for a metal. The locally averaged horizontal velocity is represented as a continuous color rainbow. In the top frame, the flyer plate (shown at the left in blue) is moving to the right and is about to collide with the sample (shown in greenish-orange moving to the left). The relative velocity of collision is 25% of the metal sound speed. The resulting shock waves move through the material, which ultimately fractures at the point where two rarefaction waves meet. Significant plastic work results in a damaged spall fragment exiting to the right.

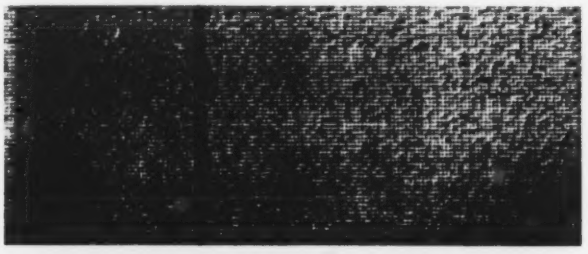

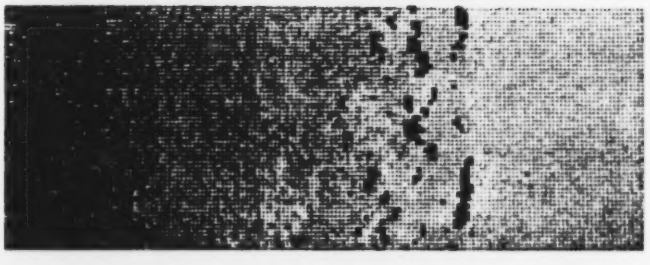

Work supported by DOE.

ince the 1986 discovery of a new class of superconducting ceramic oxides, scientists around the world have been involved in intensive research efforts to understand and fabricate practical high-temperature superconductors. These efforts to model the structural, vibrational, and electronic properties of matter can be greatly assisted by powerful computers, especially since parallel supercomputers and sophisticated algorithms can now be combined to provide sufficient processing power.

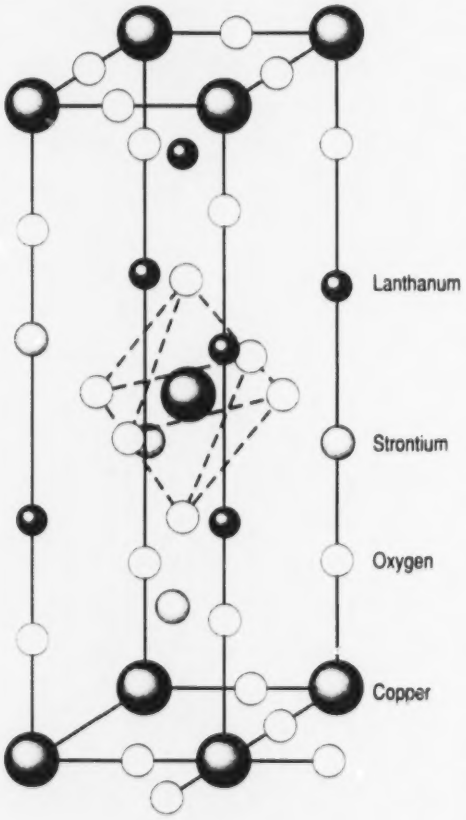

Physicists and computer scientists at Oak Ridge National Laboratory (ORNL) have collaborated to develop a multipurpose parallel computer code to calculate the electronic structure of real materials from first principles based on the Korringa, Kohn, and Rostoker coherent potential approximation (KKR-CPA) theory of magnetism and alloys. The model allows multiple atoms per unit cell and is ideally suited for situations in which substitutional disorder plays an important role, such as in high-temperature superconductors. The code executes efficiently on serial computers, on shared-memory multiprocessors such as the CRAY Y-MP, and on distributed-memory multiprocessors such as the Intel iPSC/860, and a distributed network of computers.

Initial computational experiments on the high-temperature perovskite superconductors $Ba_{1-x}K_xO_3$ and $BaPb_{1-x}Bi_xO_3$—performed on the 8-processor CRAY Y-MP at the Ohio Supercomputer Center and the 128-processor Intel iPSC/860 at ORNL—have revealed several interesting details related to alloy softening of the density of states and Fermi surface nesting. Moreover, these results revealed that experiments on more complex superconductors such as $(La_{2-x}Sr_x)CuO_4$ will be able to execute at a rate of over 2.5 Gflops on the Intel iPSC/860.

Work supported by DOE.

Plasma Physics Research

The successful development of magnetic fusion or inertial fusion as an alternative energy source requires a deep and detailed knowledge of plasma phenomena. Numerical simulation has made many contributions to the basic understanding of plasma phenomena. One can point to the discovery of nonlinear effects in space plasma physics, inertial confinement physics, and magnetically confined plasmas that came about through numerical simulation. The problems are typically complex with many competing effects acting simultaneously on different time and space scales, and are inaccessible to experimental or analytical approaches. Typically, simulations have given insights into fundamentally new physical processes that are striking in the ways that nonlinearity can bring order out of complexity. Understanding plasma phenomena depends on the ability to do simulations in three dimensions with realistic geometries and long time scales. The single most important component for continuing progress is the attainment and successful application of greater computing power. There is a present and urgent need to provide realistic plasma simulations in three physical dimensions involving complex bounded systems.

Another area with potentially important applications involves the interdisciplinary coupling of molecular physics with plasma science. For example, to be able to model the synthesis of novel materials using chemical vapor deposition (CVD), one would need to combine the computational techniques of molecular physics with numerical plasma simulations.

Fundamental Physics and QCD

Field theories such as quantum electrodynamics and quantum chromodynamics (QCD) are believed to describe the entire range of physical phenomena on the atomic, nuclear and subnuclear scales. QCD is the theory of strongly interacting particles such as nucleons and mesons. The highly nonlinear regimes of these theories cannot be solved with traditional techniques of quantum mechanics. However, an approach in which space-time is approximated by a lattice makes it possible to predict the characteristics of strongly interacting particles. The present calculations are forced to make severe approximations. With the continuing development of new computational techniques as applied to novel computer architectures it is possible to treat these more exactly. The goal of these calculations is to learn how quarks and gluons form composite

hown above is a visualization of the pion propagator based on a QCD lattice generated on the Connection Machine. The propagator is a function of three spatial dimensions and one time dimension, so the data have been averaged over the third spatial dimension and displayed as a function of x and y (the short axes) and time (the long axis).

The event represented is the creation of a pion near the center of the volume and its propagation in space both forward and backward in time. The magnitude of the propagator determines the size of the "bubbles" (shown in green) in this visualization, and a representative surface of constant amplitude is displayed in white. From the rate at which the amplitude dies out as a function of time, the pion mass can be calculated. The generation of the lattice requires approximately 300 hours on a 16K CM-2; many such lattices are required to give a statistical estimate of the pion mass by averaging results from the propagator calculation based on each lattice.

Work supported by DOE.

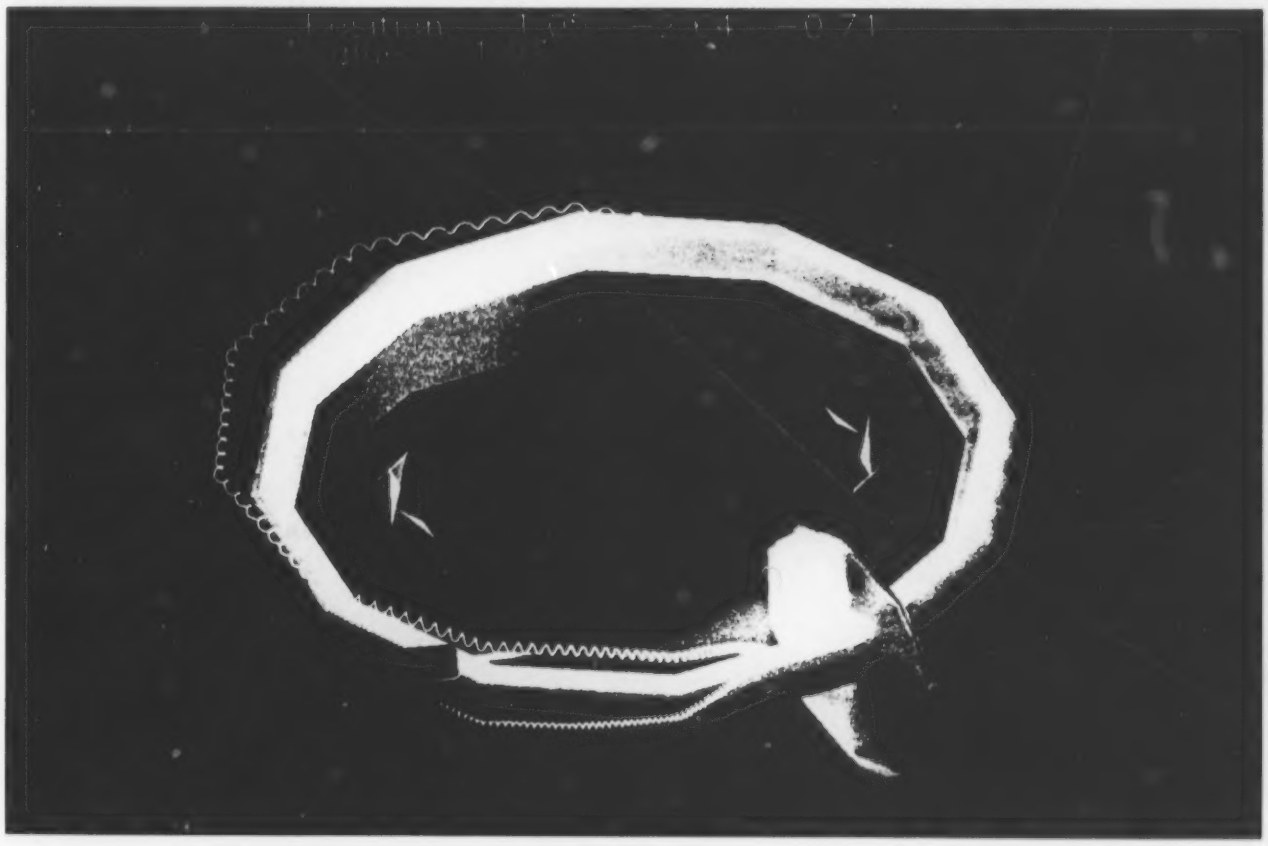

One of the more complex plasma problems to be solved is the modeling of plasma confinement by magnetic fields in a Tokamak. This modeling requires a full three-dimensional simulation of the diffusion of a multicomponent plasma in the twisted toroidal fields. The accompanying picture shows output from a three-dimensional finite difference adaptive mesh simulation. Shown is an interior surface and a cross section of the Tokamak with the color of the surface proportional to the strength of the magnetic field. The use of color allows the display of the field strength on multiple surfaces through partial transparency of one of them. The magnetic field decreases from the inner edge of the torus (yellow) to the outer edge (blue). Color shading enables one to display multiple complex features in three-dimensional computer models.

Within the Tokamak is depicted the trajectory of a single charged particle (indicated by the white line). The orbit is calculated from the magnetic field data stored on the three-dimensional mesh. The orbit generally follows a magnetic field line, while precessing slowly due to gradients in the field. This calculation is only a start on this exceedingly complex problem. Diffusion is a result of the drifting motion in this magnetic field and the self-consistent electric fields generated inside the plasma. A realistic simulation of a Tokamak will require a computation mesh with a million zones and several million particles. Edge effects, plasma heating from beams, and wave heating, as well as the internal diffusion process, must be taken into account.

Work supported by DOE.

BEST COPY AVAILABLE

structures as will be studied experimentally at Relativistic Heavy Ion Collider and Superconducting Super Collider.

QCD is among the most computer-intensive problems known. One reason is because full QCD is a very complicated field theory, with many degrees of freedom. Today's best calculations are being done with 16×16×16×32 lattices. It is thought that definitive results will not be obtained until we can calculate with 256×256×256×256 space-time sites and each site requires nearly 100 floating point words to describe the fields. That is around 400 billion words of computer memory required just to hold the description of the lattice. Furthermore, nearly all arithmetic operations in QCD involve complex arithmetic, so the number of operations per word of the lattice is high. Extracting these results will require computers thousands of times more powerful than today's fastest supercomputers.

Environmental Modeling and Remediation

A comprehensive mathematical modeling and computational research program is needed in the area of restoring the earth's surface and subsurface environment after contamination. Problems in this area are of immense importance to DOE and are found in every region in our country. Examples include leaky radioactive or chemical storage containers underground (gas stations, nuclear waste sites, etc.) and chemical spills on the surface (oil spills in the ocean, chemical processing plant spills, etc.). Computer models can help provide much needed understanding of the complex physical, chemical, and biological processes involved, and detailed simulations are effective substitutes for experimental laboratories that can be used to test understanding of complex phenomena and supplement physical intuition. In many cases computer simulations are the only feasible method of studying the phenomena due to the long time scales involved in the radioactive and chemical processes.

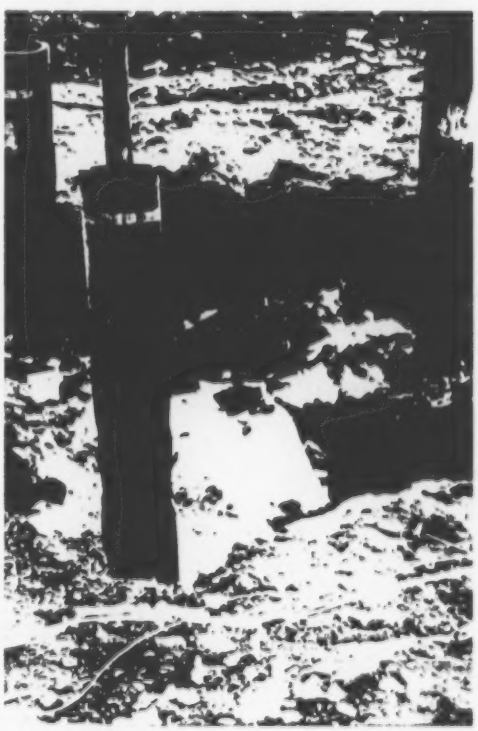

sing the heat generated by an electric field to melt a region of contaminated soil is the first step of the "in-situ vitrification" process currently under study for removing hazardous waste from DOE sites. The soil resolidifies as a glass, effectively trapping the contaminants and preventing their migration into the groundwater supply. An experiment testing the procedure has been conducted at Oak Ridge National Laboratory. The glassified soil can be seen between the four electrodes in the test configuration shown in the photograph.

Computational modeling is essential to better understand and control this process. DOE researchers are developing algorithms for modeling the many complex processes involved in the vitrification process: heat generation and transfer, natural convection, liquid/solid phase change, chemical interactions, and the creation and movement of gas bubbles. As the process takes place underground, computational models must also be used to interpret the data obtained from necessarily indirect experimental measuring techniques.

Work supported by DOE.

Computer studies are essential for the development of new technologies for identifying, monitoring, and reclaiming hazardous waste sites and spills. Many such technologies are emerging, such as *in situ* vitrification, bioremediation, organic volitalization, and surface and subsurface hydrology. Some, no doubt, will fail to live up to their potential, but others may provide solutions that significantly improve the process. Even a small improvement can yield tremendous savings given the magnitude of the problem facing our environment today. Computer simulations will help to determine which technologies should be vigorously pursued and under what conditions they should be utilized.

Current models incorporating relatively few chemical, biological, and physical interactions run at 20 to 40 Mflops on a single CRAY X-MP processor and three-year simulations take on the order of 10 hours of CPU time and over a day of actual wall clock time. When additional interactions, such as inclusion of more important chemicals in the reactions, cracks in subsurface rock formations, and the potential exponential growth of biological species, are incorporated along the required resolution to resolve the chemical reactions at the fronts, computational requirements quickly surpass the capabilities of today's supercomputers, even for small three-year simulations. Because of the time scales involved, simulations of ten to a few hundred years are required, which further enlarge the computational power shortage.

SOFTWARE COMPONENTS AND TOOLS

In addition to a need for software tailored specifically for grand challenges, one of the HPCC Program objectives is to advance generic software technology that will have broad national impact. DOE plans to form collaborative groups among the national laboratories, universities, and industry to develop and share in this software technology effort. This area includes a broad range of activities, a selection of which are described below.

Massively Parallel Software Systems

Massively parallel computers will not displace traditional supercomputers for the majority of users until reliable, multiuser operating systems are developed. Such operating systems are needed to ensure efficient machine utilization and should be developed along with the hardware to ensure adequate hardware support for operating system functions. The operating systems that are currently available for parallel computers are still primitive compared to UNIX and MACH. Although a multitasking, timesharing system might eventually be desirable, the initial operating systems may assign each user a subset of the total number of nodes on a massively parallel machine. Using this approach would avoid problems associated with load balancing while timesharing occurs on some fraction of the nodes. The operating system must also ensure the integrity and security of each application being executed on the computer and allow operators to restart all of the applications after an inadvertent shutdown. A very important issue is the ability of an operating system to allow input/output to occur in parallel. Without such a capability, an input/output bottleneck can form, limiting the utilization of the machine. It is important to capitalize on recent advances in a universal parallel, distributed-memory operating system such as parallel MACH, which will have a small kernel and a large library of servers containing most of the UNIX functionality. This will allow streamlining the operating system, providing only those services on individual nodes needed by a given job. The operating system must provide functionality, tying kernels on individual processors together. This should include parallel file systems that can be smoothly created and accessed from several cooperating processors. The advantages to both manufacturers and users of a universal operating system are obvious: leverage of resources and uniformity of environments.

Portable and Scalable Libraries

Mathematical software has played a vital role in making uniprocessors effectively usable by scientists and engineers. As more elaborate computer architectures arise, the need for more sophisticated mathematical software becomes acute. Ideally, such mathematical libraries should be portable across a variety of architectures, and scalable libraries for a single architecture will be the first step in this evolution. To ensure portability across machines, one may adopt a language for expressing vector and parallel constructs, such as an extended version of Fortran. It is clear, however, that such a common language alone cannot guarantee "performance portability." Until compilers for parallel machines with vector or RISC processors mature, avoiding substantial degradation across architectures will still rely on (1) tuning code by inserting compiler directives, and (2) using basic primitives that are written in assembler for maximum exploitation of architectural features and that are a part of the mathematical library of a given machine. To ensure performance

portability (and avoid degradation of performance), one needs numerical libraries in which the algorithms make effective use of the primitives that ideally match the architectural feature for maximum performance. For example, scalability of the algorithms used in these libraries can be enhanced by scalable primitives. This has been demonstrated, to a certain extent, in libraries that deal with dense matrix computations. However, much remains to be done in other areas, such as sparse matrix computations.

Programming Languages/Compilers

Effective use of high performance computers depends heavily on both the programmer's ability to express algorithms in a programming language clearly and succinctly and the associated compiler's ability to exploit the target machine's architecture. There is extensive scientific debate on the most cost-effective language/compiler strategy for parallel architectures. The principle options are (1) extend existing languages with parallelism primitives, (2) advance compiler optimization technology to find parallelism in existing languages, and (3) develop functional, object-oriented languages to replace the old styles. The DOE laboratories will collaborate with interested university and industrial partners to use their extensive expertise to push the technology of all three options. They will also establish fair comparisons among the options to determine their relative merit.

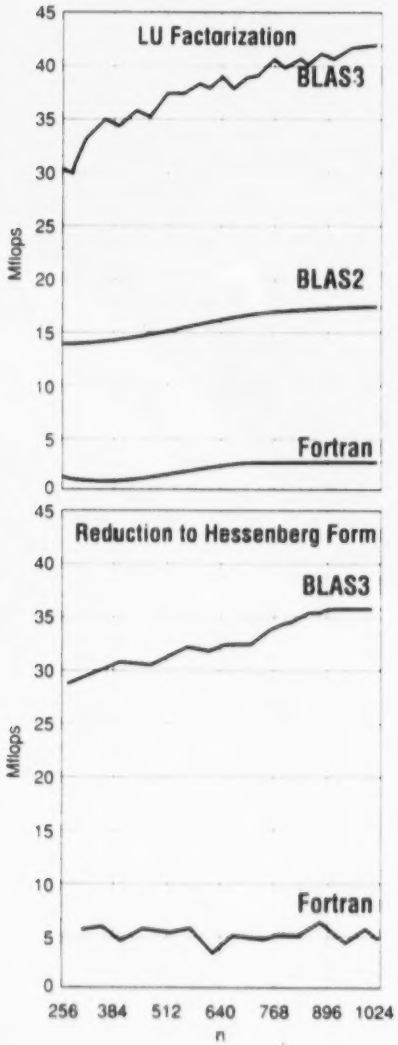

evelopment of Basic Linear Algebra Subroutines (BLAS3) was based on the idea that the addressing locality of a program can be greatly enhanced by properly indexing a set of nested loops. For example, indexing through whole rows and columns can cause many more cache misses than repeatedly indexing through small rectangular blocks that fit in cache memory. In the case of the Alliant FX-8, cache memory misses caused performance to fall far short of expectation. With BLAS3, performance increased to almost peak achievable system performance, which is sustained for problems that are not cache contained. Performance rates of approximately 50 megaflops have been observed for the rank-k update BLAS3 primitive C←C+B. The performance of the block methods that use these primitives is also quite high. Rates between 40 and 45 megaflops have been observed for a block LU decomposition algorithm.

Top: The performance of an LU factorization implemented in Fortran, using optimized BLAS2 primitives and BLAS3 primitives.

Bottom: The results of the BLAS3-based code for reducing a matrix to upper Hessenberg form and the Eispack Fortran routine ORTHES after automatic optimization.

Work supported by DOE, NSF, AFOSR, and IBM.

oad balance is a critical issue in porting applications to massively parallel computers. At DOE's Sandia National Laboratories, a dynamic synchronous load balance method based on binary decomposition was used to balance one million particles in six seconds on a 1024-processor NCUBE/ten starting from a random assignment of particles to processors; subsequent load balance updates required less than 0.1 second each. Using heterogeneous programming in a radar simulation application, six user programs—a host program, dynamic scheduler, ray-tracer, radar simulator, image collector, and graphics program—execute simultaneously and cooperatively on the 1024-processor NCUBE 2. The dynamic scheduler performed asynchronous load balancing based on a processor hierarchy; the resulting code ran about 50 times faster on the NCUBE than on a CRAY Y-MP processor. This scheme has been recently applied to SDI tracking problems and other complex simulations.

Work supported by DOE.

Scientists at Argonne National Laboratory played a leading part in the design of Strand, a parallel programming system distributed by Strand Software Technologies of Beaverton, Oregon. Recently, members of Argonne's Mathematics and Computer Science Division helped the Aerospace Corporation of El Segundo, California apply Strand to engineering problems. They also helped the corporation develop a set of tools for understanding the performance of parallel programs.

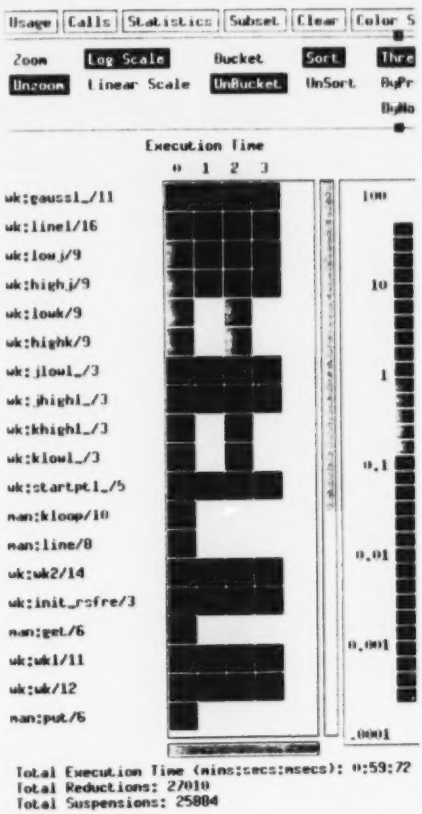

The parallel processing group at the Aerospace Corporation originally contacted Argonne because they needed parallel programming tools to support parallel versions of codes used in their engineering programs. They were also interested in extending tools that they had developed for understanding the performance of sequential programs to work in a parallel programming system.

Argonne scientists helped Aerospace Corporation develop their tools into a form suitable for parallel performance evaluation. The result of this work is a powerful data collection and analysis system called Gauge. Gauge can be used on a wide variety of parallel computers, including hypercubes and shared-memory machines. Engineers at Aerospace Corporation are now using Strand and Gauge to develop parallel versions of target tracking and rocket design codes. At Argonne, the Gauge tools have proved invaluable in projects developing parallel codes of climate modeling and computational biology.

The group at Argonne also advised Strand Software Technologies on how to incorporate the Gauge tool into the Strand programming system. Strand Software Technologies agreed to distribute Gauge free of charge with the Strand software system, providing wide distribution of the tools developed by the collaboration.

The different colors in the image represent the amount of work being performed in different parts of the parallel program. Black represents the most work, whereas red, blue, and green represent less work, in that order (see color bar).

Work supported by DOE.

treams and Iterations in a Single Assignment Language (Sisal) is a general-purpose functional language for high-performance computing. The language is intended to run on both conventional and novel multiprocessor systems. The language's development is a collaborative effort by Lawrence Livermore National Laboratory and Colorado State University. Currently, researchers at more than twenty academic and research institutions worldwide are using Sisal.

Functional languages provide a clean and easy-to-use parallel programming model that facilitates parallel program development and simplifies compilation. The semantics of functional languages isolate the user from the complexities of parallel programming. The compiler and operating system—not the user—are responsible for the synchronization, communications, and scheduling of concurrent tasks. Functional languages free the user to concentrate on the nature of the algorithm and not its execution.

Until now, functional semantics carried a high performance cost; but recent advances in compiler and operating system technology have brought Sisal on par with Fortran on both scalar and vector shared-memory multiprocessors. Given the expressive and easy-to-use parallel programming model it provides, Sisal is an attractive alternative to conventional programming languages on shared-memory multi-processors.

Program	# Lines	1 Processor Sisal	1 Processor Fortran	5 Processors Sisal	5 Processors Fortran
Gauss Jordan	114	2.6	1.7	0.9	1.0
RICARD	301	7.3	7.7	1.8	1.8
SIMPLE	1150	128.7	128.8	32.9	99.6
Weather Code	2718	55.5	75.1	13.6	54.4

This table, a comparison of Sisal and Fortran execution times, shows the execution times (in seconds) of four scientific programs on the Alliant FX/80. Clearly, Sisal is as fast as or faster than Fortran.

Work supported by DOE and ARO.

esearchers at the University of Illinois' Center for Supercomputing Research and Development have pioneered compiler techniques to detect implicit parallelism in sequential languages, such as Fortran 77. Today, these compiler techniques are used routinely in most supercomputers. For algorithm kernels, the parallel code produced by restructuring compilers is comparable to that generated manually. The goal of much current research is a similar achievement for large applications.

Work supported by DOE.

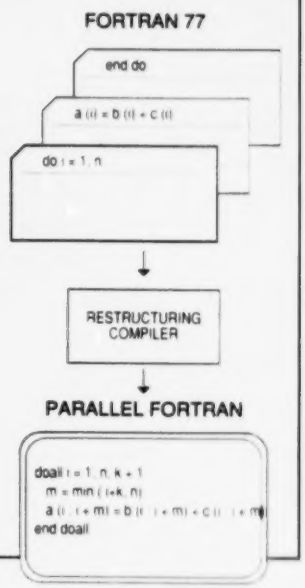

Computational Science Environments

Much remains to be accomplished in order to create effective work environments for computational scientists. The issues range from programming tools to facilitate error-free code generation, to performance analysis and enhancement, from pre- and postprocessing of data, to effective visualization of multidimensional time-dependent phenomena, and from communication of scientific results to the creation of problem-solving environments oriented toward particular classes of applications and using the specialized idioms and language of those applications areas. The objective is to make HPCC and the computational science approach to problem solving more useful for and readily available to large numbers of engineers and scientists. This is particularly relevant to the transfer of computational science technology to industry.

Tools that merit early consideration include common, multimedia user interfaces, graphic job monitoring and management aids, parallel Fortran and C extensions, parallel multinode debuggers, performance analyzers, shared data cache systems, transport, hierachical file storage systems checkpoint/restart software, and automatic memory coherency managers.

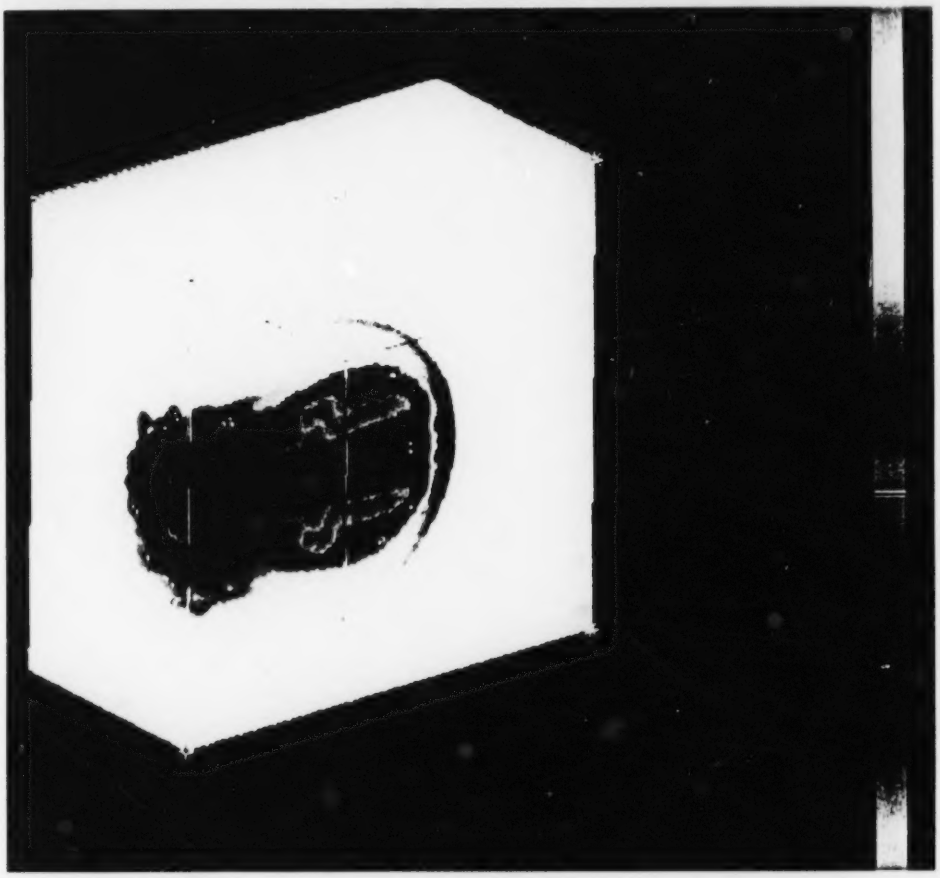

os Alamos National Laboratory has developed MediaView, a generic framework for communication via multimedia documents. These documents can include not only text, line art, and still images, but sound, video sequences, and computer-produced animations. Also, when cast in digital form, the mathematical content of a document can be symbolically and numerically manipulated. Thus, one can experiment with the mathematics, derive new results, and simulate different situations with different parameters. Finally, an animated scientific visualization can be incorporated in a MediaView document and be examined in situ *or electronically mailed to a colleague for independent study. Thus, MediaView is a communication tool that offers new and dramatically different ways of interacting with others.*

Media View is easy to use and understand. It is based on a word processor metaphor, something familiar to any computer user. In addition to text, that metaphor is extended to include several multimedia components. And like text, these additional components are subject to the same cut, copy, and paste paradigm, making them as simple to manipulate as words. As a result, powerful and complex MediaView documents can be constructed by nonspecialists.

MediaView can be of enormous benefit to workers in a distributed but networked HPC "center." One of many possible examples of use is exchanging scientific visualizations and mathematical analyses. Unlike a conventional document, these components appearing in MediaView are live and interactive.

The figure is an interactive animator for a sequence of images that were produced on a supercomputer. The colors represent the density of the fluid calculation being simulated here. The color bar, from cyan to yellow, indicates less dense to more dense. Different colors delineate structures that are formed in the fluid simulation. The observer is viewing a cut-away of a Mach-6 intergalactic jet simulation, with the green lines showing the boundary of the data, had a "slice" not been cut-away.

Work supported by DOE.

BEST COPY

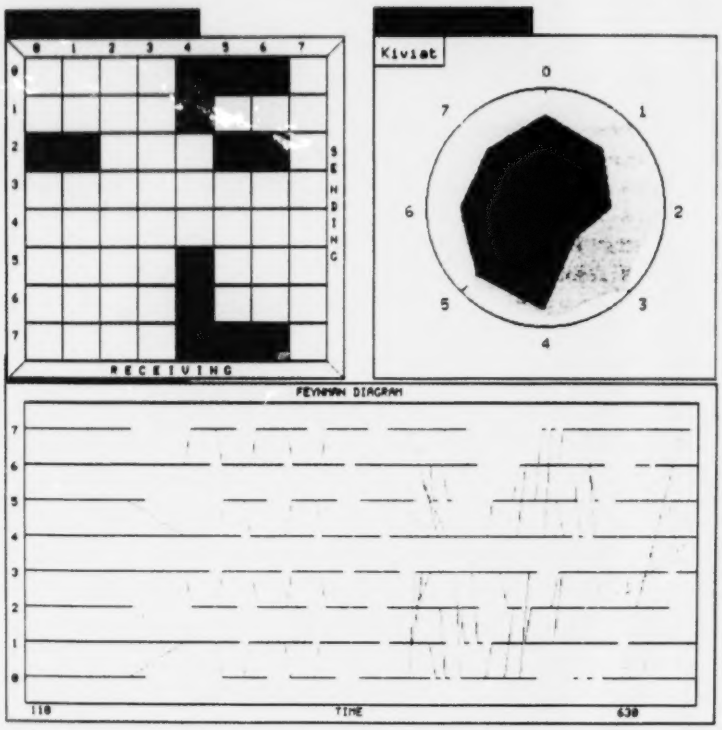

esearch and development at Oak Ridge National Laboratory in the area of performance characterization for parallel supercomputers has produced two software tools for monitoring and visualizing the behavior of parallel algorithms. Portable Instrumented Communication Library (PICL) is a subroutine library that implements a generic message-passing interface on a variety of multiprocessors. Programs that use PICL routines for interprocessor communication are portable in the sense that the same source code can be compiled and executed on many different parallel architectures. Optionally, PICL also provides time-stamped trace data on interprocessor communication, processor busy/idle times and user-defined events. These data serve as input to a second tool, called ParaGraph, which is a graphical display system for visualizing the behavior of parallel algorithms on message-passing multiprocessor architectures. ParaGraph uses both color and motion to provide a dynamic visual depiction of the behavior of the parallel program. ParaGraph provides several distinct visual perspectives from which to view the same performance data, in an attempt to gain insights that might be missed by any single view.

Three of the approximately twenty views in ParaGraph are shown in the accompanying figure. The circular image (two shades of red) in the upper-right portion of the diagram shows how efficiently the program is running on the processors and the load balance between processors. The larger, lighter-shaded polygon indicates the maximum efficiency of the program thus far for each processor, whereas the smaller, darker ploygon shows the current efficiency and load balance. The larger the polygon, the more efficient the algorithm, and the more circular, the better the load is balanced. Ideally, one would want the darker polygon to always be large and circular and about the same size as the lighter polygon.

Work supported by DOE.

At the Georgia Institute of Technology, a research group led by physicist Uzi Landman has been a major user of the Department's supercomputer facilities for a number of years. Landman's specialty is the computerized simulation of the atomic world, particularly as it is reflected in the properties of materials and in the physical and technological consequences of their microscopic interactions. By incorporating both classical and quantum mechanical principles into his supercomputer programs, Landman has been able to accurately predict how various materials behave—at the molecular and atomic level—under various physical conditions and to explain how they respond when brought into contact with each other under pressure. Visualization of these simulations is providing scientists with fundamental new knowledge of the atomic mechanisms that underlie the behavior of materials. Landman, shown in the photo, illustrates on a monitor how an electrically unbalanced cluster of salt molecules (Na^{14+} Cl^{13+}) becomes stabilized. To make up for the missing chlorine atom, a free electron (indicated by the cloud of white dots) attaches itself to a sodium atom (shown in red) on the surface of the cluster and begins to pull it away, leaving a balanced cluster of 13 sodium and 13 chlorine atoms.

Work supported by DOE.

Distributed Computing

Computational science is a diverse and geographically distributed activity. As previously mentioned and as discussed further below, one of the goals of the HPCC Program is to establish HPC Research Centers that will be used by many researchers via high-speed networks. Many researchers will also have significant local resources in the form of powerful workstations and local shared fast microprocessor systems. Effective utilization of these distributed DOE resources in attacking grand challenge problems will require software to support distributed systems. Distributed systems are systems in which two or more computers cooperate on the solution of the problem. For example, one may want to distribute computations in order to match algorithms for specific subproblems with the most appropriate computer architecture. The goal of distributed systems research is to create a uniform user and programming environment across heterogeneous systems; in other words, to make it as easy to use a collection of computers on a given problem as it is to use a single computer. Important areas of research and development include distributed programming environments to make it easy to create, debug, and run distributed scientific applications; distributed operating system techniques for control of distributed resources; distributed file system management and archival support tools to facilitate utilization of high performance networks; and mechanisms for effective distributed software control and update.

Photo courtesy RAY STANYARD

Uzi Landman

Visualization and Imaging

The purpose of HPCC technology is to augment the work of scientists and engineers. Their interface to these systems is of central importance if they are going to be most effective. Therefore, it is important to provide the hardware and software technology to allow for input of a variety of image data and for output that supports the use of their senses, particularly visual, for gaining maximum problem insight.

Research and development is needed in input, storage, and analysis of image data and in the output of data in appropriate visual and other sensory forms. The output techniques needed include effective management of three-

Supercomputers do not operate as stand-alone devices. Typically they are embedded in a heterogeneous shared resource environment consisting of many services, for example: mass storage, input/output, authentication, graphics engines, terminals, and workstations. During the 1970s Lawrence Livermore National Laboratory (LLNL) played a leading role in developing the shared resource server model of a supercomputer environment. During the 1980s LLNL developed a distributed operating system architecture for a supercomputer environment. The goal of this research and development was to create the programming and user-at-a-terminal logical model of a single shared and nonshared memory multiprocessor system providing transparent resource access. The architecture was implemented as the native client-server, multiprocessing operating system on CRAY X-MP and Y-MP systems that has been in production for over two years. Several network services were implemented within the architecture. For example, the storage part of the architecture was implemented on UNIX systems as a distributable hierarchical mass storage system. LLNL is collaborating with General Atomics to make this storage system available to industry and other supercomputer centers. The architecture also provided the basis for the IEEE Distributed Mass Storage Reference Model.

Work supported by DOE.

dimensional data and display, browsing, connection of output images to quantitative data, and comparison of current and past or alternative approaches. Visualization and imaging are dependent on high performance networks, effective data-compression techniques, scientific databases, and high-definition recording and display technology.

Very Large Scientific Databases

Many of the grand challenge applications of interest to the DOE—such as global climate modeling, understanding the structure of matter, and human genome—potentially involve very large scientific databases containing hundreds or perhaps thousands of terabytes of data. These databases will require hierarchies of storage media, which in turn will require new approaches to database structuring, storage, and searching. In addition, improved interfaces are needed to facilitate scientists' querying, browsing, studying, and transforming the databases to enable better understanding of the physical phenomena portrayed in them. Standards need to be developed in this area to aid collaboration among scientists and the aggregation of logically single databases from distributed databases being generated at many experimental locations.

COMPUTATIONAL TECHNIQUES

The principal objective in this research area will be to improve the performance of numerical simulations on parallel computing systems. Research in this area should not only encompass algorithm kernels and subroutines, but should also include full-scale scientific and engineering applications of contemporary interest.

Computational Methods and Algorithms

DOE will support a comprehensive research program in advanced computational techniques. The principal goal of this effort will be to provide methodologies and algorithms to enable the effective solution of grand challenge problems and to promote U.S. industrial competitiveness. The focus will be on methods that allow exploitation of the protypical architectures deployed at the HPC Research Centers. Creation of the high performance algorithms, to be incorporated in the methodology, will require, in many instances, new numerical techniques that exploit greater parallelism than can be simply extracted from parallelized versions of existing sequential algorithms.

Applications Algorithm Design Systems

In addition to developing new methods and algorithms, it is important to package these techniques and make them easily usable. An attractive area for research is how to create parallel design methodologies for large-scale scientific applications that include modular and well-integrated software tools, user interfaces, standards and protocols. Such design methodologies will expose the similarities and important distinctions between the hardware architectures and provide mechanisms to promote portability. New fundamental mathematical algorithms will also be posed in such frameworks.

HIGH PERFORMANCE COMPUTING RESEARCH CENTERS

DOE will contribute to the support of HPC Research Centers. The main purpose of these centers will be to enable computational scientists to explore the effectiveness of new full-scale high performance architectures in solving grand challenge problems. By using full-scale architectures, code developers can carry out experimental code design, with the expectation of adequate compute power to solve the problem. In order for researchers at remote locations to make effective use of these facilities, each center will become a major node on a National Research and Education Network backbone. Collaborative consortia of DOE-supported HPC Research Centers, industry (both user industries and computing and communications industries), and universities will be very highly encouraged.

The sponsorship of the HPC Research Centers is an area in which DOE has a long history of excellence and leadership. DOE national laboratories have recognized the importance of computational experiments and analysis to complement expensive physical experiments. Many of the DOE national laboratories continue to operate major production-oriented supercomputing centers, as well as a number of parallel computing research centers. There are many types of parallel computer architectures that will need evaluation. Because several already have displayed high potential at outperforming contemporary vector multiprocessors, it is reasonable to expect that there will be additional such architectures identified in the High Performance Computing Systems component of the HPCC Program. Among the concepts to be thoroughly evaluated are SIMD, MIMD, and hybrid processors and shared, distributed, and hierarchical memories. Each hardware architec-

he development of a comprehensive and general information system for molecular biologists—the Chromosome Information System (CIS)—is a major part of the computing effort at the Lawrence Berkeley Laboratory (LBL) Human Genome Center. In this system, one consistent graphical user interface will transparently interact with multiple underlying databases. The separation of user actions on the biological data objects and the data storage enhances the system flexibility and functionality. The description of the data, and the operations on the data, are specified in a high-level language that captures the biological concepts while permitting the use of commercial data management systems at the lower levels. Through integration of specialized databases (e.g., an image database, a sequence database, and a map database) biological information may be accessed at many levels of abstraction.

This work is based on concepts developed by the Data Management research group at LBL. The CIS has been implemented by a collaboration of the Data Management group and biologists from the LBL Human Genome Center.

Work supported by DOE.

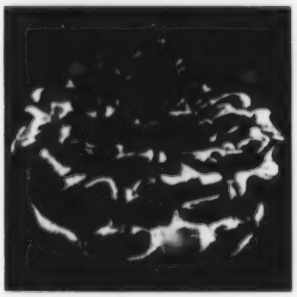

Efficient modeling of three-dimensional, time-dependent fluid flows plays a critical role in many grand challenge research problems. Adequate resolution of these flows using conventional spatial discretizations would require computational effort far in excess of the most optimistic projections for near-term growth in computational capabilities. Scientists at Lawrence Livermore National Laboratory, the Courant Institute, and Los Alamos National Laboratory are developing adaptive methods that focus computational effort where it is most needed, resulting in a dramatic increase in the computational efficiency and enabling the solution of realistic 3D flow problems.

One type of adaptive technique is local adaptive mesh refinement, in which the computational mesh dynamically changes as a function of space, time, and data to maintain a fixed level of accuracy in the calculation. The figure shows computational results obtained using this method to model the collapse of an ellipsoidal cloud of Freon when it is hit by a shock wave, where yellow indicates the location and concentration of the Freon. The use of adaptive refinement for this problem reduced computational costs by more than an order of magnitude, resulting in a savings of hundreds of Cray hours.

Work supported by DOE, DARPA, NSF, AFOSR, and DNA.

ture type will be evaluated with respect to competing software architectures for the various grand challenge applications.

The activities at each center will be designed to support interdisciplinary and inter-institutional collaborations. A key ingredient will be a critical mass of in-house research conducted by focused teams of applications scientists, computational mathematicians, and computer scientists. These collaborative teams will direct their efforts toward the solution of grand challenge problems of interest to DOE. They will be responsible for maintaining a dialogue with universities, industry, and other laboratories and centers in order to maximize the dissemination of information and avoid unnecessary duplication of effort.

The HPC Research Centers will also play an important role in computational science education. This aspect of their mission is expected to go well beyond the usual training courses in computing and to embrace such issues as computational science curriculum design at the university level and introductory and motivational programs and material for high school and elementary students and their teachers. Clearly, accomplishing these activities will require strong collaboration with the education community.

During FY 1985, the U.S. Congress mandated a study of computer networking requirements of the U.S. research community. The preparation of such a far reaching study, together with the growing widespread remote access to supercomputers through federal supercomputer programs, such as those in the DOE and NSF, focused the attention of the U.S. research community on the need for more capable, high-speed computer networks.

Within the DOE, computer networks had already been used extensively for specific applications and programs, primarily supporting Fusion and High Energy Physics, but these networks were mostly incompatible and lacking in capacity. Because of this and because of a recognized and significant increase in networking requirements, the Energy Research (ER) community endorsed a proposal to create the Energy Sciences Network (ESNet). The ESNet has been developed to be compatible with existing network requirements while providing connectivity and interoperability to other federal research networks in addition to a nondisruptive transition path to emerging international network standards.

The ESNet is the vehicle through which the ER community has become a full partner in the Internet community of computer networks and through which the ER community will become an integral part of the proposed National Research and Education Network (NREN). In fact, the ESNet currently incorporates access by

SNet is currently a T1-based (1.5 Mbit/s) data communications network supporting more than twenty major Office of Energy Research (OER) sites directly connected to the backbone. It is a multiprotocol network supporting two major families of data communications prototcols, DECnet and the Internet protocol.

The five major OER programs supported are Basic Energy Sciences, Health and Environmental Research, High Energy and Nuclear Physics, Fusion Energy, and the Superconducting SuperCollider.

ESNet is engineered, installed, and operated by the networking staff of the National Energy Research Supercomputer Center (NERSC), located at Lawrence Livermore National Laboratory. Initial deployment of the T1 circuits began in late 1989 and became fully operational with both protocol families by early 1990. Since becoming operational, total traffic on the network has doubled about every six months.

In addition to interconnecting the major OER-supported sites (shown on the map) ESNet provides access to several regional networks, has international connections to Japan and Europe, and currently connects to two Federal Inter-agency eXchange (FIX) points providing interconnects to other national backbone networks, including the NASA Science Network and the National Science Foundation Network (NSFnet).

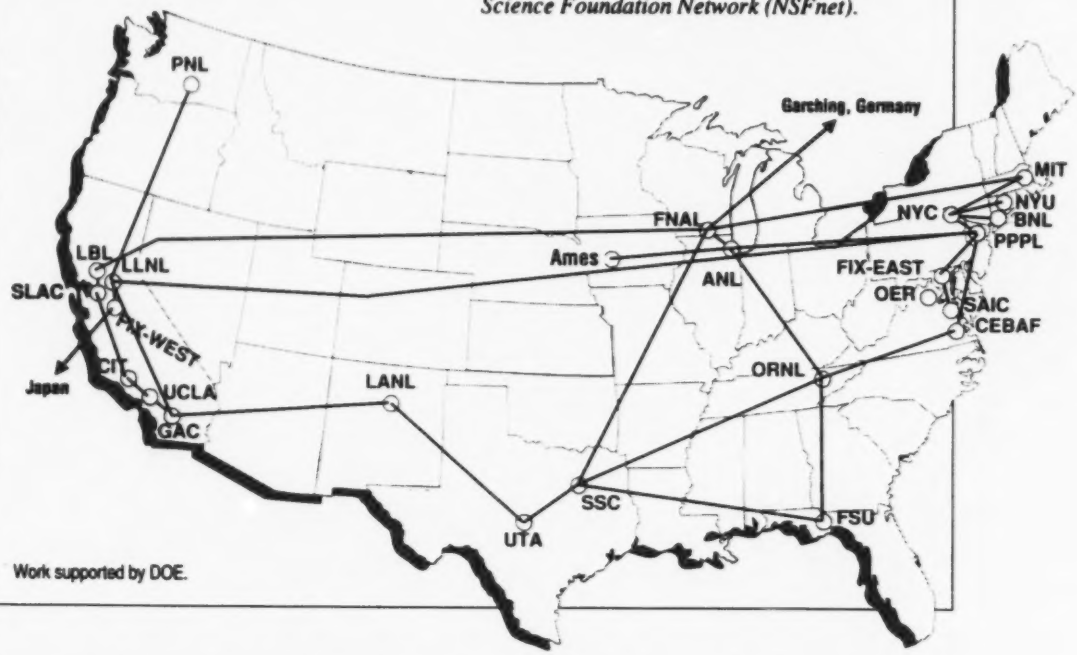

many ER collaborators through the NSFnet and through regional networks as was proposed in the ESNet Program Plan, June 1987.

These networks, along with ESNet and others, are generally considered to be the precursor to the NREN.

The ESNet has been proposed by the DOE within the framework of the HPCC Program, of which NREN is a major component. The HPCC Program is a balanced computing research proposal that includes research in HPCC systems, software technology, and computer networks as well as funding for human resources and the NREN. The NREN, as proposed, will be a computer communications network that interconnects educational institutions; national laboratories; non-profit institutions; government facilities; commercial organizations engaged in government supported research; and unique national scientific and scholarly resources such as supercomputer centers, libraries, etc.

The NREN will provide high-speed communications access to over 1300 institutions across the U.S. within the initial planning period, and is proposed to offer sufficient capacity, performance, and functionality so that the physical distance between institutions is no longer a barrier to effective collaboration. The NREN will support access to HPCC facilities

The ESNet was developed after careful definition and documentation of user requirements from all of the Energy Research program areas. This user requirements orientation led to an ESNet implementation that departed from other Internet implementations. Most notably, the ESNet was developed to be a multiprotocol network to preserve Energy Research user applications investments in existing network implementations while allowing full interoperability with other Internet implementations, such as the NASA Science Internet and the NSFnet. Recently, the Internet Advisory Board has adopted the concept of incorporating a multiprotocol approach to allow a nondisruptive transition to international standards.

Work supported by DOE.

and services such as full-motion video, rapid transfer of high-resolution images, real-time display of time-dependent graphics, remote operation of experiments, and advanced information sharing and exchange. The NREN is also intended to incorporate advanced security and a uniform network interface to domestic users as well as a standard interface to international research networks.

The NREN proposes to achieve economics of scale by serving many federal agencies, industrial R&D centers, and university campuses. Although the federal government will provide a substantial direct investment in the NREN, it is important to recognize that a large indirect investment will be made by academic and industrial institutions and other networks that will connect to the NREN. The DOE NREN budget is for the direct support of DOE laboratory and other DOE-funded research facility connections to the NREN, per the existing federal budget process. This DOE program will complement other NREN deployment, funded and coordinated by the NSF, to the broader research and education community.

The ESNet is user requirement driven within the framework of the DOE mission. The DOE HPCC Program proposal for the NREN is applications and user oriented, both with regard to the operations and the gigabit research areas. The DOE will participate in the NREN management to ensure a requirements-driven approach with strong user involvement and site coordination. Because of the common mission orientation and common requirements of the ER community, the consensus of this community is that the ESNet would continue to be viewed as a large, logical community of interest network within the NREN context, even if in the longer term ESNet were no longer a separate but compatible physical network.

Several of the most demanding functional requirements for DOE networking capabilities are highlighted below.

The ER community will have a continuing need for multiprotocol support throughout the initial planning period. Although there is a constant transition to UNIX-like systems and TCP/IP network protocol usage, a wholesale conversion to these systems is impractical and expensive and cannot be accomplished within three years. During that time, OSI standard protocols will have been incorporated into the ESNet, where they will co-exist with TCP/IP, X.25, and DECnet.

Videoconferencing has been initiated within the ESNet in a test mode. In recent years, the cost of video technology, the use of compression techniques for video transmission, and the use of packetized video have spurred great interest in using video for research collaboration support. The ESNet video pilot project has received very positive evaluations from all involved.

The benefits of video are even more evident as one examines the logistics of the many international collaborations in which ER programs are involved. Video enables these collaborators and negotiators to meet on a weekly and even daily basis, which could not be done otherwise.

The use of windowing technologies and other software tools has enabled more facile and productive use of remote computing and control facilities. These tools 'extend' the capability of a principal investigator's local workstation to these remote facilities in such a fashion that they can be used more readily in the investigation of complex problems. The operation of these tools in a distributed environment, however, is a very taxing network requirement.

In addition to windowing tools, there is a growing requirement for distributed programming and debugging tools. Remote procedure call functionality has recently been implemented on the supercomputer systems at the National Energy Research Supercomputer Center, and as this functionality achieves more widespread use, it is anticipated that the distributed processes will place a heavy demand on ESNet bandwidths. It will be very important to have the ability to remotely analyze system performance with regard to parallelization, individual system node processes, and throughput, etc., especially, when evaluating prototype 100-gigaflops and teraflops systems as part of the HPCC Program. This will place even more demands on network bandwidth (possibly another order of magnitude) in the FY 1993-95 timeframe.

Proposals for the next large experiments within the Fusion Program for the International Thermonuclear Experimental Reactor and within the High-Energy Physics Program for the Superconducting Super Collider include the requirement for distributed control, i.e., remote operation, of the experiment. Presently, such requirements for experiment control require the most capable local area network implementations. These requirements for the wide area ESNet, in the FY 1995-96 timeframe, will require bandwidths in the gigabit range.

DOE NETWORKING RESEARCH PLAN

The DOE HPCC networking plan must encompass at least the following technical research issues to meet anticipated near term needs for using ESNet and NREN as they evolve.

Performance

Virtually all the characterization to date of the technical parameters of NREN has centered around bandwidth. However, additional factors such as response time and availability can be equally, or even more, important for many applications and must be given corresponding consideration. In addition, support will be provided for research in network performance evaluation.

Policy-Based Routing

Policy-based routing is the term used to denote the routing of traffic, taking into account the usage policies of both the using entity and the service provider. This issue is still in an early research phase. Topics of study include determining mechanisms and parameters to allow independent domains to protect themselves from unwanted traffic and mechanisms to determine paths through the system meeting client needs.

Open Systems Interconnect (OSI) Integration

The OSI networking standards need to be incorporated into the NREN architecture and capabilities.

Multiprotocol Support

A variety of protocol suites are currently used by various user communities. For DOE, these include DOD IP, DECnet (Phase 4), SNA, X.25, and others. User-level requirements based on these protocols will continue for some time. The network must allow for support of these protocols if all or most users are to be supported in their research. Design of gateway and transition mechanisms between protocol suites will require research and development. New higher performance protocols will also be investigated.

Network Management and Operations

One of the most important factors in providing gigibit network service across a large community will be the ability to establish a hierarchical network management and operations enterprise that will be able to effectively deal with network management issues in a very diverse environment.

Problem Resolution Tools

User level problems in networking are becoming increasingly sophisticated and subtle, while at the same time the complexity of the network environment is rapidly increasing. Resolution of these problems will require a sophisticated arsenal of tools that will help to confirm, analyze, and pinpoint problems at all levels of network services.

PRIVACY AND SECURITY

Protecting the NREN and its attached resources against attack is an important requirement. Protection and monitoring mechanisms must be part of its design.

The above is a partial list of considerations that must be addressed by the NREN, in addition to the major charter of providing state-of-the-art performance to the HPCC Program. An interagency NREN will evolve over time, adding these capabilities as they can be reliably provided. It is unlikely that these divergent requirements can be met with other than a hierarchical architecture, incorporating multiple networking entities and multiple levels of networking components, many of which may be in place already.

GIGABITS RESEARCH AND DEVELOPMENT

The leadership for gigabit network research for the Federal HPCC Program will be performed by DARPA. DOE will support a modest effort in gigabit network research: to support grand challenge applications, to provide connectivity within and among its HPC Centers, to contribute the high performance networking expertise of the DOE laboratories to the HPCC Program generally, and to ensure that it maintains the expertise to upgrade and integrate its ESNet with the emerging gigabit NREN. The gigabit network research program below meets these needs. The previous section outlined key research issues needed for the interagency interim NREN effort; they apply equally to the full NREN effort. Below are listed additional gigabit research and development efforts that DOE will pursue.

Local Gigabit Networks

For a national gigabit backbone to be fully effective, local gigabit networks are required to interconnect the local resources of the HPC Centers. Two main approaches to local gigabit interconnection currently exist: the circuit switch approach and the bus approach. For each, there are many unanswered questions

The CASA wide-area testbed is one of five very high speed communication network projects lead by the Corporation for National Research Initiatives (CNRI). CASA is truly an interagency collaboration involving the California Institute of Technology, DOE's Los Alamos National Laboratory, NASA's Jet Propulsion Laboratory, and NSF's San Diego Supercomputer Center. MCI, Pacific Bell, and U.S. West are also collaborating with the CASA testbed.

Primary research goal of this collaboration is the effective simulation of grand challenge problems on geographically dispersed supercomputing resources connected via very high speed networks. The communications environment envisioned for CASA relies on the Los Alamos-proposed High Performance Parallel Interface (HIPPI) ANSI standard and will leverage research at Los Alamos on HIPPI-based crossbar switches, networking software, and protocol processors. The CASA project will develop several prototype distributed supercomputing applications in chemistry, geophysics, and global climate modeling. The present supercomputing resources available in the collaboration are Thinking Machines Corporation's CM-2, Cray Research Incorporated's Y-MPs, and several hypercube architectures. As with the other testbeds, the CASA effort will contribute to the research and development base required for the National Research and Education Network.

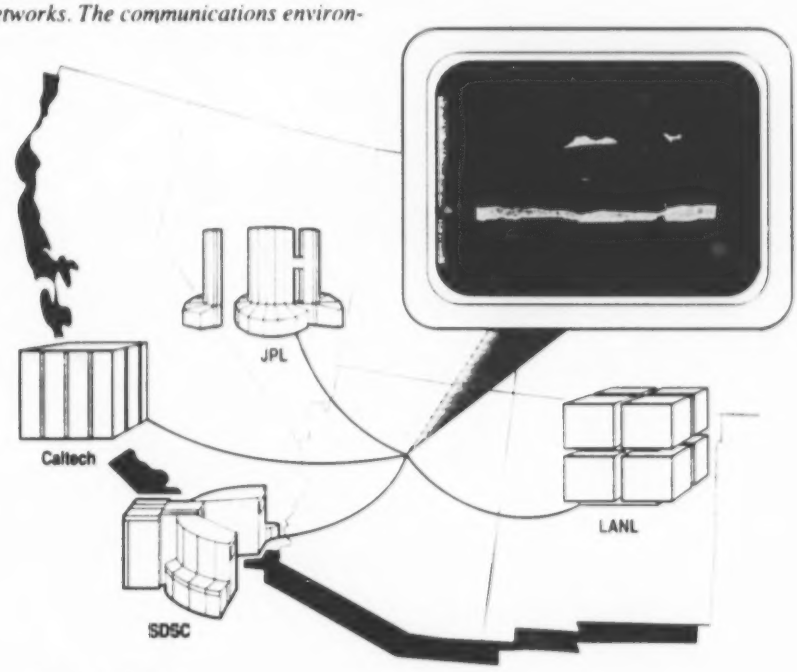

Work supported by NSF and DARPA.

about the best design. The circuit switch approach looks particularly attractive because of its excellent bandwidth scaling, security, and connection setup overlap characteristics. DOE, using existing gigabit circuit switch expertise and equipment within its national laboratories, will investigate circuit switch architectural questions, such as how is the number of connections economically scaled, where is the optimum placement of functionality, and what is the effect of various host and network interface I/O architectures on actual throughput.

Gigabit Applications

Prototype local and long-haul gigabit network applications need to be developed that demonstrate the utility of these networks for enhancing the scientists' productivity, achieving computational power through distributed resource sharing, or enabling other, possibly unforeseen, significant scientific or economic advantages for the grand challenge applications. Developing distributed applications, using resources of several HPC Centers cooperatively on a single problem, may require

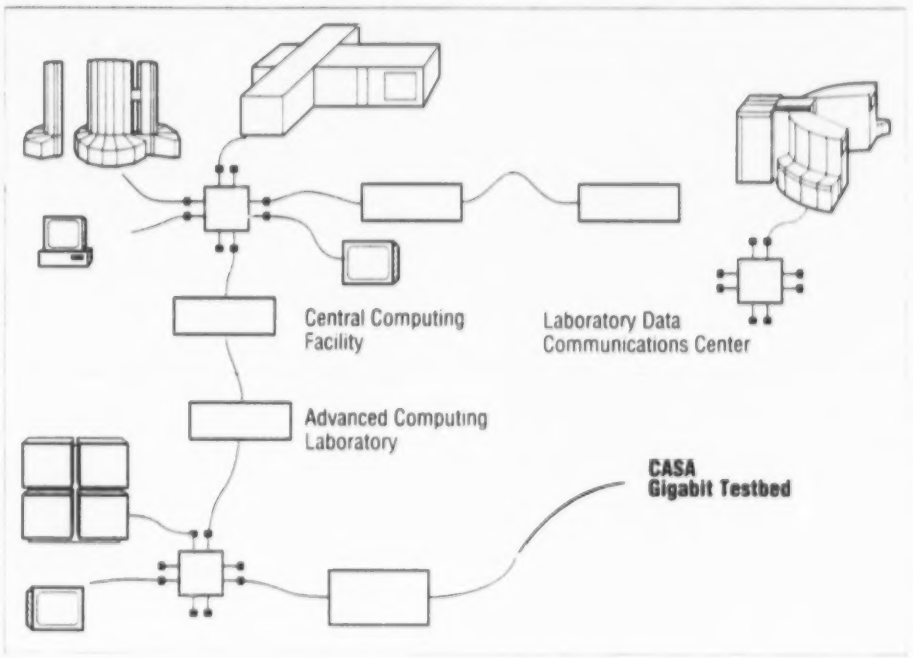

he Multiple Crossbar Network (MCN) testbed is a gigabit per second testbed network at Los Alamos National Laboratory that interconnects computing resources in the Advanced Computing Laboratory, Central Computing Facility, and Laboratory Data and Communications Center. It will be extended in approximately a year to include the wide-area CASA gigabit per second testbed. CASA sites include Caltech, the Jet Propulsion Laboratory and the San Diego Supercomputer Center.

The purpose of the MCN testbed is to experiment with gigabit per second networking technology, to understand how to apply this technology to leading-edge scientific problems, and to promote closer interaction with industry. The local distribution system for the MCN testbed is based on work done in the Laboratory's Network Engineering Group on the High-Performance Parallel Interface (HIPPI) and HIPPI-based switching systems. A Network Systems Corporation crossbar switch with a 6.4 gigabit per second aggregate throughput and less than 1 microsecond switching time is being coupled to Los Alamos-designed intelligent network interfaces to assemble a high-performance HIPPI switching system called CP. Several CP* systems will be interconnected with fiber-optic links to assemble a multiple crossbar network.*

This network will interconnect supercomputers such as CRAY X-MPs and Y-MPs, Thinking Machine CM-2, workstations from Sun and IBM, framebuffers, and high-performance disk systems. This extensive collection of systems with HIPPI interfaces provides a unique facility to do performance and interoperability testing of their HIPPI implementations.

The applications for this testbed include visualization of high-resolution (1024 × 1024 × 24) images at video rates (24 frames/second), distributed supercomputing using software tools such as Express and ISIS, and data motion. This testbed will provide facilities to explore the applications and networking requirements for high-performance computing systems of the 1990s.

Work supported by DOE.

To improve network congestion control, the Network Research group at Lawrence Berkeley Laboratory developed algorithms for slow-start, dynamic window allocation and improved round-trip estimators. These improvements enhanced performance over the Internet (loaded) by factors of 2 to 100 and the technology was adapted by virtually all major computer manufacturers.

The group also developed a header compression algorithm to enable practical implementations of serial line IP protocol, which permits computers to act as network hosts over modems, thereby providing simpler, more reliable, and more secure data connections. This technology, too, has been transferred to industry.

Other research includes development of a TCP header prediction algorithm to increase throughput on high-speed networks, investigation of policy-based routing, and development of a new algorithm that allows network addresses to have arbitrarily complex structure yet permits routing look-ups to be completed in constant time. This algorithm—the only one known to handle ISO addresses efficiently—will be included in 4.4 BSD UNIX.

Work supported by DOE.

new mathematical algorithms, new application structuring, and new interprocess communication techniques and protocols.

Protocols and System Software

With high performance computers and gigabit local area networks in place, research will be performed to measure communication performance and to identify bottlenecks in applications, operating systems, and communication protocol architectures and implementations. These studies will lead to experiments with new approaches to operating system and protocol design and implementation. Extensions of parallel computing tools to enable applications programs to run on multiple hosts linked with gigabit wide-area networks will be investigated.

Gigabit Testbeds

Because it is important to experiment with alternative gigabit networking approaches, local gigabit testbeds will be set up associated with the HPC Centers. DARPA and NSF have already initiated a significant Gigabit Testbed Research Program. DOE-supported sites participate in these national or regional area testbeds. DOE will provide additional testbed resources between sites for networking application experimentation that does not interfere with production networks.

BASIC RESEARCH AND HUMAN RESOURCES

This component of DOE's HPCC Program will build on the existing DOE applied mathematical sciences base program and recognizes that a portion of the HPCC funding must go toward long-range investment in the future. While the other three components are designed for a near- or intermediate-term impact, basic research funding must increase for the long-term health of the field. Likewise, to develop the human resource potential of the future, investment in education must accelerate, beginning at the high school level or below and continuing through life-long training. In both basic research and education, DOE has a long history of contributions and experience from which to draw. Our plan for the HPCC Program uses this base and builds upon it.

BASIC RESEARCH

Basic research supports all other aspects of the HPCC Program. The topics below address fundamental aspects of computing and its applications.

New Algorithms

Part of the vitality of HPCC depends on a steady influx of new algorithms and mathematical models. Indeed, the current generation of practical, reliable algorithms, which are the tools for attacking the grand challenges were once exotic new techniques designed and tested on model problems. DOE has long provided a natural and effective platform for research into new algorithms. Techniques that have grown out of DOE's basic research program now serve as general tools for hundreds of different applications. The continued renewal of creative algorithms and mathematical models is one of DOE's primary strengths, and the HPCC Program will exploit this strength.

Parallel Numerical Algorithms

Much of HPCC is scientific numerical modeling. As the sophistication of the physics models increases, the complexity of numerical problems expands. Parallel computers magnify the issue. We need continuous development of the best parallel numerical methods to ensure that we can truly apply the power of our most advanced computers to model the needed physics correctly.

Matching Algorithms to Architectures

Research on HPCC has produced a diverse set of multiprocessor designs. Each design seems to favor certain types of algorithms and perform poorly on others. The relationship between good matches is often not obvious at the outset. We need to have a better understanding of how to best select and use a particular machine for solving a spectrum of different computational problems.

Performance Measurement and Analysis

When a parallel program runs poorly on a specific machine, the cause of the slowness can be baffling. It could be the algorithm, the translator, the partitioning of work, communication bottlenecks, system overheads, hardware failures, or any number of other causes. Moreover, many current analysis techniques alter the performance as they try to measure it. Mechanisms for accurate and timely identification of the cause of poor performance must be improved. Successful use of large numbers of processors will depend on progress here.

Novel Applications

Many important applications of HPCC remain to be discovered and explored. Those who best understand the potential of HPCC must ensure that a process for continually exploring new applications is developed and implemented. When feasibility has been demonstrated and documented for a novel application, responsibility for its continued exploitation will transfer to the affected applications discipline.

Theory of Parallel Computing Complexity

Traditional complexity work provides a foundation for identifying efficient sequential algorithms, but it does not directly apply to parallel machines. Parallel complexity models often make unrealistic simplifications, such as an infinite number of processors or an infinite supply of memory, an assumption that leads to many impractical algorithms. We need a solid foundation for parallel computing algorithms that will accurately represent their performance on large, but finite, numbers of processors and various types of parallel computers.

Software Productivity and Reliability

The cost of software for conventional computers remains very high. The software cost for parallel systems will be substantially greater. Likewise, software reliability is often questionable. We need to develop a practical methodology for building high-quality software that can

In recent years, DOE researchers have developed and analyzed advanced numerical methods for solving the large-scale nonlinear differential and algebraic systems that arise during the numerical solution of ordinary and partial differential equations. This research has applications in a diverse range of problem areas, including circuit simulation, computational electromagnetics, computer-aided design of mechanical systems, modeling of chemically reacting flow, power systems, chemical vapor deposition and optimal control. The development of effective solution techniques has required a new understanding of the mathematical structure of differential equation systems, as well as a rethinking of algorithms. For extremely large systems, new techniques for large-scale linear and nonlinear systems have been developed.

New algorithms and methodologies have been transferred to industry in part through the development of high-quality mathematical software. Software developed by DOE for the solution of these problems now has hundreds of users in industry, government and academia; variants of this software have been incorporated into several commercial codes. The analysis and software have directly impacted solution techniques for the dynamic analysis of mechanical systems, trajectory control, and for a wide variety of chemical engineering applications. Second-generation algorithms and software are designed to attack extremely large systems. Effectiveness of these techniques was recently demonstrated on a 2D laser oscillator model at Lawrence Livermore National Laboratory with 38,000 unknowns. Work is now under way to exploit the excellent potential for parallelization of these techniques.

Work supported by DOE, ARO, AFOSR, NASF, Aerospace Corp., and Exxon.

erformance variations in modern supercomputers can be visualized by plotting the megaflops delivered on the Perfect Benchmarks. Each dot in the figure represents that megaflops rate of the indicated machine on one of the Perfect codes. The gaps between theoretical peak rates and actual rates delivered on real applications demonstrate the need for comprehensive performance studies of the relationships between supercomputer architectures and applications and systems software.

Work supported by DOE, NSF, and NASA.

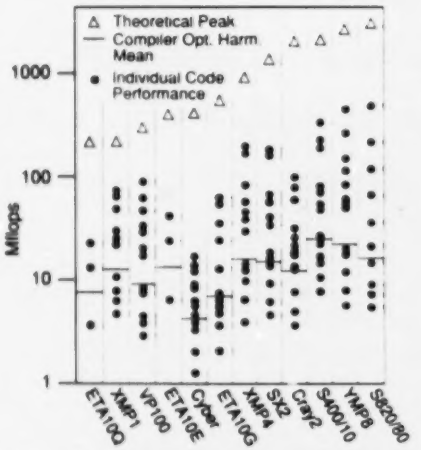

ost computer benchmarks are simple, fixed-size problems that are timed for various machines. Ames Laboratory researchers have recently completed the design of the first computer benchmark based on fixed-time comparison principles. Called SLALOM, for Scalable, Language-independent, Ames Laboratory One-minute Measurement, the benchmark adjusts the number of unknown variables in a complete radiation transfer problem such that the input, setup, solution, and output times take a total of one minute. It answers the question, "How big a problem can one run with this computer?" It is the first benchmark to allow fair comparison of the whole spectrum of computer speeds and types, from the smallest personal computers to the fastest vector and parallel supercomputers. SLALOM is rapidly gaining acceptance from computer manufacturers as a just and realistic performance metric that has been needed for some time.

Work supported by DOE.

esearch on software tools at Argonne National Laboratory has been directed at the problem of aiding users who are adapting existing scientific and engineering computer programs to execute efficiently on parallel machines. The research is motivated by the enormous investments in expertise and computing resources that such programs represent. Because the work of a community of users may depend on the effective application of a program, its adaptation to parallel execution has significant scientific and economic consequences. A collection of prototype tools that emerged from this research has been further developed into a commercial software product named VecPar_77 by the Numerical Algorithms Group, Inc. (NAG), which will market the product worldwide.

VecPar_77 is a pre-compiler tool that analyzes and transforms Fortran programs. Its function is illustrated by its application to the parallelization of a computer program—developed by scientists at Argonne, NASA, and IBM—that uses multigrid methods to model fluid flow in two dimensions. In the existing serial Fortran code, four types of operations are carried out by sweeping sequentially over rectangular subdomains called patches. *Parallelizing the code becomes a matter of carrying out these operations concurrently over the patches. VecPar_77 extracted information about dependencies that inhibit parallel execution of the key loops that carry out these operations. Armed with this information, the program developers made changes that enables these loops to be executed in parallel. The result is a significant performance improvement on machines that permit such parallel operation—e.g., the Sequent Symmetry.*

The commercial development of VecPar_77 was a cooperative endeavor in which one of the Argonne scientists who had participated in the research took leave to act as a consultant to NAG. His efforts were partially supported by a technology commercialization grant from the State of Illinois, administered by the Technology Transfer Center at Argonne.

Work supported by DOE and the state of Illinois.

cientists at Argonne National Laboratory are determining the feasibility of using logic programming on high-performance computers to study two- and three-dimensional structure in molecular biology.

A group at Argonne has devised a computer program that exploits logic programming techniques to match up to 500 sequences of DNA/RNA (basic genetic building blocks of organisms). It is the first such automated code to handle more than a modest number (about ten) of sequences. The program, which is based on a combination of the logic programming language Prolog and a lower-level language C, runs efficiently on the Sequent Symmetry machine and a network of Sun workstations.

The Argonne group, in collaboration with scientists at the University of Illinois in Urbana, has also developed a computer-based procedure that accurately computes the structure of ribosomal RNA. Ribosomal RNA forms the core of the translation apparatus—the mechanism that reads genetic information into protein sequences. It thus lies at the heart of the cell's function and replication.

Argonne's automated procedure has already proved successful in several ways. Previously undetected secondary structure has been identified at positions 63 and 104 of E. coli *16S ribosomal RNA (shown on the diagram in red). Moreover, new tertiary structure has been discovered in 7S RNA.*

Since the chemical activity of ribosomes is involved in a number of processes—such as building and repairing tissues, digestion, and the immune system's response to disease—understanding the structure of ribosomal RNA is of considerable scientific and medical interest.

Work supported by DOE.

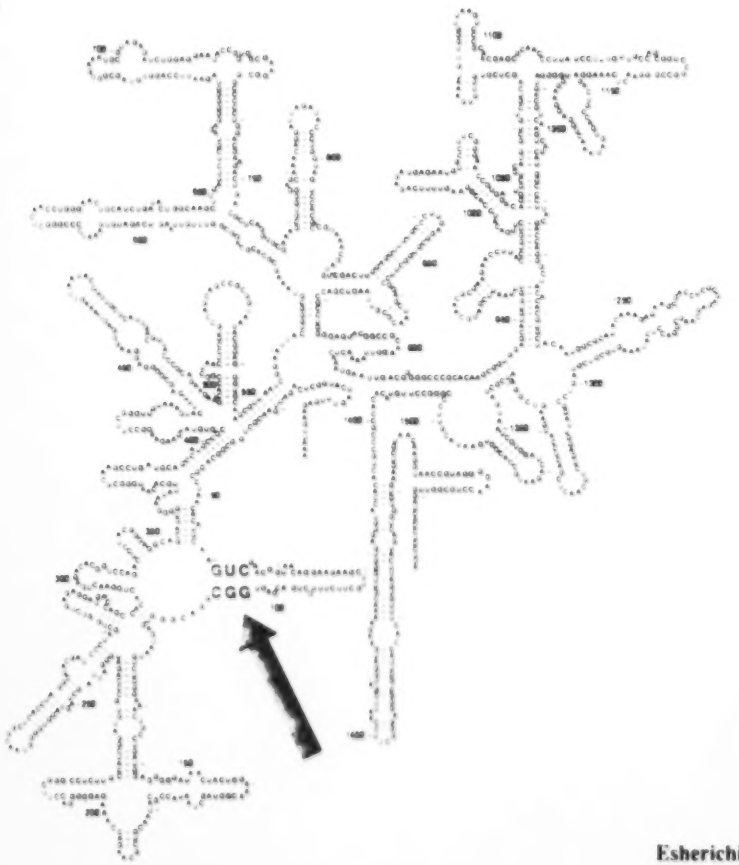

Esherichia coli

be used and expanded by others and with capabilities that are clearly and accurately delimited as an integral part of the code.

Parallel Reasoning Systems

HPCC offers tremendous potential for nonnumerical applications. We anticipate that future-generation supercomputers will actually function as scientific assistants, performing routine tasks with great efficiency and speed and making informed decisions on the basis of accumulated experience. We need to explore the introduction of parallel programming paradigms into automated reasoning and logic programming.

HUMAN RESOURCES

The long-term viability of U.S. competitiveness requires a workforce highly educated in engineering and scientific disciplines. In recent years traditional theoretical and experimental techniques in science and engineering have been augmented by a powerful new technique: computational science. Scientists and engineers now use computers for basic research in complex physical, biological, and chemical phenomena, as well as in the design and manufacturing of complex engineering systems including automobiles, airplanes, nuclear reactors, and computers.

While theory and experimental methods will not be replaced by computational science, they are being augmented by it in important ways. Computer simulations bridge theory and experiment by simultaneously considering sophisticated models and controlled study of complex results. Computing also glues scientific disciplines together. Sophisticated tools are required to solve the basic equations on a computer, independent of the particular application. For example, the equations used to approximate wing flow also describe gas flow in automobile engines, blood flow in the heart, waves in the ocean, hydraulic machinery, fluids in computer disk drives, flow of polymers in material science, motion of oil in the earth, spread of pollutants in the atmosphere, microchip fabrication, and the formation of metal alloys, as well as many other complex phenomena. A large body of expertise is common to these seemingly different applications.

Currently, computational science programs are emerging in a small number of U.S. universities. DOE laboratories have long been leaders in developing computational science techniques and in training limited numbers of postdoctoral researchers. However, the number of postdoctoral visitors at the national laboratories in mathematics, computational science, and any form of engineering is very low relative to other areas.

In the future, high school science and mathematics courses will introduce computational science techniques, students will use them during college, and they will continue to develop them throughout their professional lives. DOE will play a substantial role in supporting our educational system and in helping to provide the highly specialized training beyond that which is possible within traditional degree programs. The following are ways that DOE plans to contribute:

High School Programs to Attract Talented Students

DOE will significantly expand its current summer program for inviting high school students and high school teachers to visit national laboratories. In addition, it will establish close ties with high school teachers to assist in appropriate curriculum development, to provide guest lecturers, and to enhance their use of computers in the classroom, including access to supercomputers. This type of early exposure to computing should draw more students into computational science. In addition, it will generate a U.S. workforce that better understands how computers can improve their productivity, regardless of their ultimate career choice.

Undergraduate Programs to Augment Educational Training

DOE laboratories currently have limited summer and co-op programs for college students that provide an opportunity for students to see how their formal training applies in practice. These programs also provide valuable feedback and support to schools on the status of their programs. DOE plans to enlarge these programs and extend them to top students from a greater number of colleges and universities. The emerging university graduate programs in computational science will need support in the next decade to evolve undergraduate programs. This support will include access to advanced computers for classroom teaching, summer jobs in computational science for undergraduates, and development of course materials that cut across traditional departments.

Graduate Programs to Focus Students on Research Opportunities

At the graduate level, we need to interest the best students in solving tough, important problems. DOE can contribute to this process in several ways, including jointly funded research projects; participation of DOE researchers on Ph.D. thesis committees; sponsorship of research and curriculum development workshops; creation of pre- and postdoctoral fellowships for computational science in universities; and establishment of sabbatical opportunities for faculty and their students. Many of these activities take place now; additional DOE funding through the HPCC Program will permit them to take place on a much larger scale.

Post-Graduate Programs to Advance Research

New Ph.D.s have the skill and knowledge to tackle the enabling technology needs and grand challenges of the HPCC Program. DOE has started a postdoctoral program at the labs to strengthen research efforts in key areas. For the HPCC Program, it is expected that the number of DOE-funded postdoctoral positions will increase substantially at the national laboratories and at universities.

University/Industry Exchanges to Expand Technology Transfer

Technology transfer addresses those people already in the workforce. DOE must provide a wide range of training to expand staff's skill in the rapidly changing area of computational science. DOE can also use its existing expertise by encouraging temporary personnel exchanges. DOE staff could help build new programs in computational science within high schools and universities while exposing faculty to the nature and needs of the national laboratories. Likewise, exchanges with industry could quickly accelerate the process of taking technology developed at the laboratories and put it to use in industry. In the past, political, social, and legislative obstacles have severely limited this type of activity; with the HPCC Program, DOE hopes to change this situation.

he Computational Science Graduate Fellowship Program, sponsored by the DOE's Office of Energy Research, is designed to support highly capable science and engineering students interested in pursuing doctoral study in an applied science or an engineering discipline with applications in high performance computing. The program provides a stipend ($18,000 the first year, $19,200 the second year, $20,400 the third year, and $21,600 the fourth year), tuition and fees, an institutional allowance ($1,000 a year), some travel expenses, and funding for a workstation.

The program is open to U.S. citizens who have a bachelor's degree in life or physical sciences, engineering, or mathematics. Eligible applicants may be entering graduate school or have extensive graduate study. Students who have received department (faculty) approval of a Ph.D. thesis topic are not eligible for the program.

Universities must apply and be approved for participation in the program. Acceptance will be based on material submitted in the university application. These materials include a description of the curriculum, enrollment data, previous and ongoing research, postgraduate employment record, and faculty résumés.

Student applications consist of undergraduate and graduate transcripts (GRE scores will be required after the 1991–1992 cycle), faculty references, an academic and career goal statement, and a list of work experiences and publications.

All fellows are required to work at a DOE or a DOE-approved facility in a research assignment related to ongoing high performance computing activities.

For application forms and more information, contact

Computational Science Graduate
Fellowship Program
Science/Engineering Education Division
Oak Ridge Associated Universities
P.O. Box 117, Oak Ridge, TN 37830-0117
(615) 576-0128; Telefax (615) 576-0202

Work supported by DOE.

CONCLUSION

This document describes the DOE program component of the Federal High Performance Computing and Communications (HPCC) Program. The HPCC Program features increased cooperation between industry, academia, and government and defines a strategy for supporting advances through directed R&D efforts, reduced uncertainties to industry for development and use of new technologies, support for underlying research, network, and computational infrastructures, and support for the U.S. human resource base to meet the needs of industry, academia, and government.

The HPCC Program consists of four complementary components in key areas of HPCC. These areas are HPCC systems, advanced software technology and algorithms, the National Research and Education Network, and basic research and human resources. The HPCC Program is driven by the need for unprecedented computational power to investigate and understand a wide range of scientific and engineering problems that are referred to as grand challenges and that are the fundamental problems of investigation for the mission of DOE and the other participating agencies.

The DOE has a long history of reliance on and excellence in high performance computing and computational science. Consequently, DOE fully embraces the goals of the HPCC Program and intends to mobilize its scientific community to help accomplish them.

Several of the themes that will pervade DOE's HPCC Program are already clear; others will emerge as the program is implemented and evolves. The program will be applications driven but highly interdisciplinary, involving

extensive teaming of scientists and engineers with mathematicians and computer scientists. Collaboration with industry and universities will be strongly encouraged. HPC Research Centers will provide a focus for HPCC Program intellectual activities. Computational science education will be emphasized to insure future human resources for the program.

DOE will make use of the field work proposal system to fund HPCC activities at the national laboratories. Unsolicited proposals submitted under the Special Research Grant Program (see DOE/ER-0249, October 1985) will be used to fund HPCC Program activities at other organizations.

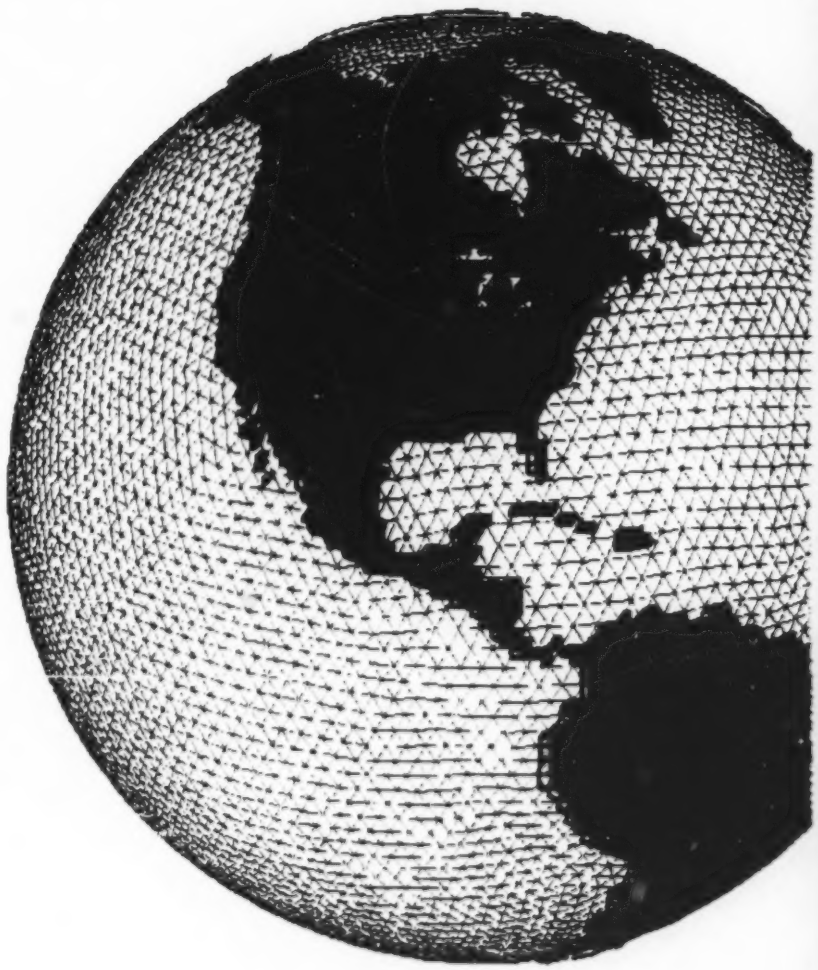

REFERENCES

1. *Grand Challenges: High Performance Computing and Communications*, to obtain a copy of this document, send request to Committee of Physcial, Mathematical, and Engineering Sciences, c/o National Science Foundation, Computer and Information Science and Engineering, 1800 G. Street NW, Washington DC 20550.
2. *Building an Advanced Climate Model: Program Plan for the CHAMMP Climate Modeling Program*, U.S. Department of Energy report DOE/ER-0479T, December 1990.

BEST COPY AVAILABLE

Prepared for the
Department of Energy's
Office of Energy Research
Scientific Computing Staff

Most Informational Publications from the Department of Energy are printed in one color. However, in order to retain technical content, four colors have been used on a few figures in this publication.

BEST COPY AVAILABLE